Query Stability in Monotonic Data-Aware Business Processes [Extended Version]*

Ognjen Savković, Elisa Marengo, and Werner Nutt

Free University of Bozen-Bolzano
39100 Bolzano, Italy
{ognjen.savkovic,elisa.marengo,werner.nutt}@unibz.it

Abstract

Organizations continuously accumulate data, often according to some business processes. If one poses a query over such data for decision support, it is important to know whether the query is *stable,* that is, whether the answers will stay the same or may change in the future because business processes may add further data. We investigate query stability for conjunctive queries. To this end, we define a formalism that combines an explicit representation of the control flow of a process with a specification of how data is read and inserted into the database. We consider different restrictions of the process model and the state of the system, such as negation in conditions, cyclic executions, read access to written data, presence of pending process instances, and the possibility to start fresh process instances. We identify for which facet combinations stability of conjunctive queries is decidable and provide encodings into variants of Datalog that are optimal with respect to the worst-case complexity of the problem.

1998 ACM Subject Classification H.2.4 [Systems]: Relational databases

Keywords and phrases Business Processes, Query Stability

arXiv:1512.06912v1 [cs.DB] 21 Dec 2015

1 Introduction

Data quality focuses on understanding how much data is fit for its intended use. This problem has been investigated in database theory, considering aspects such as consistency, currency, and completeness [8, 13, 23]. A question that these approaches consider only marginally is where data originates and how it evolves.

Although in general a database may evolve in arbitrary ways, often data are generated according to some business process, implemented in an information system that accesses the DB. We believe that analyzing how business processes generate data allows one to gather additional information on their fitness for use. In this work, we focus on a particular aspect of data quality, that is the problem whether a business process that reads from and writes into a database can affect the answer of a query or whether the answer will not change as a result of the process. We refer to this problem as query stability.

For example, consider a student registration process at a university. The university maintains a relation *Active*(*course*) with all active courses and a table *Registered*(*student*, *course*) that records which students have been registered for which course. Suppose we have a process model that does not allow processes to write into *Active* and which states that before a student is registered for a course, there must be a check that the course is active. Consider the query Q_{agro} that asks for all students registered for the MSc in Agronomics (*mscAgro*). If *mscAgro* does not occur in *Active*, then no student can be registered and the

* This report is the extended version of a paper accepted at the 19th International Conference on Database Theory (ICDT 2016), March 15-18, 2016 - Bordeaux, France.

© Ognjen Savkovic and Elisa Marengo and Werner Nutt;
licensed under Creative Commons License CC-BY

query is stable. Consider next the query $Q_{courses}$ that asks for all courses for which some student is registered. If for each active course there is at least one student registered, then again the query is stable, otherwise, it is not stable because some student could register for a so far empty active course.

In general, query results can be affected by the activities of processes in several ways. Processes may store data from outside in the database, e.g., the application details submitted by students are stored in the database. Processes may not proceed because data does not satisfy a required condition, e.g., an applicant cannot register because his degree is not among the recognized degrees. Processes may copy data from one part of a database to another one, e.g., students who passed all exams are automatically registered for the next year. Processes may interact with each other in that one process writes data that is read by another one, e.g., the grades of entry exams stored by the student office are used by academic admission committees. Finally, some activities depend on deadlines so that data cannot change before or after a deadline.

Approach Assessing query stability by leveraging on processes gives rise to several research questions. (1) What is a good model to represent processes, data and the interplay among the two? (2) How can one reason on query stability in such a model and how feasible is that? (3) What characteristics of the model may complicate reasoning?

(1) **Monotonic Data-Aware Business Process Model.** Business processes are often specified in standardized languages, such as BPMN [22], and organizations rely on engines that can run those processes (e.g., Bonita [7], Bizagi [16]). However, in these systems how the data is manipulated by the process is implicit in the code. Current theory approaches either focus on *process* modeling, representing the data in a limited way (like in Petri Nets [18]), or adopt a *data* perspective, leaving the representation of the process flow implicit [4, 6, 11]. We introduce a formalism called Monotonic Data-aware Business Processes (MDBPs). In MDBPs the process is represented as a graph. The interactions with an underlying database are expressed by annotating the graph with information on which data is read from the database and which is written into it. In MDBPs it is possible that several process instances execute the process. New information (fresh data) can be brought into the process by starting a fresh process instance (Section 2). MDBPs are monotonic in that data can only be inserted, but not deleted or updated.

(2) **Datalog Encodings.** Existing approaches aim at the verification of general (e.g. temporal) properties, for which reasoning is typically intractable [4, 10, 11]. In contrast, we study a specific property, namely stability of conjunctive queries (Section 3), over processes that only insert data. This allows us to map the problem to the one of query answering in Datalog. The encoding generates all maximal representative extensions of the database that can be produced in the process executions and checks if any new query answer is produced. We prove that our approach is optimal w.r.t. worst case complexity in the size of the data, query, process model and in the size of the entire input.

(3) **MDBP Variants.** When modeling processes and data, checking properties often becomes highly complex or undecidable. While other approaches in database theory aim at exploring the frontiers of decidability by restricting the possibility to introduce fresh data, we adopt a more bottom-up approach and focus on a simpler problem that can be approached by established database techniques. To understand the sources of complexity of our reasoning problem, we identify *five restrictions* of MDBPs: *(i)* negation is (is not) allowed in process conditions; *(ii)* the process can (cannot) start with pending instances; *(iii)* a process can (cannot) have cycles; *(iv)* a process can (cannot) read from relations that it can write; *(v)* new instances can (cannot) start at any moment. Combinations of these

restrictions define different variants of MDBPs, for which we investigate the stability problem (Sections 3–9).

Related work and conclusions end the paper (Sections 10, 11). A technical report, with complete encodings and proofs can be found in [24].

A preliminary version of this paper was presented at the AMW workshop [21].

2 Monotonic Data-Aware Business Processes

Monotonic Data-aware Business Processes (MDBPs) are the formalism by which we represent business processes and the way they manipulate data. We rely on this formalism to perform reasoning on query stability.

Notation We adopt standard notation from databases. In particular, we assume an infinite set of relation symbols, an infinite set of constants *dom* as the *domain of values*, and the positive rationals $\mathbb{Q}^+$ as the *domain of timestamps*. A schema is a finite set of relation symbols. A *database instance* is a finite set of ground atoms, called *facts*, over a schema and the *domain* $dom_{\mathbb{Q}^+} = dom \cup \mathbb{Q}^+$. We use upper-case letters for variables, lower-case for constants, and overline for tuples, e.g., $\bar{c}$.

An MDBP is a pair $\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$, consisting of a *process model* $\mathcal{P}$ and a *configuration* $\mathcal{C}$. The process model defines how and under which conditions actions change data stored in the configuration. The configuration is dynamic, consisting of *(i)* a database, and *(ii)* the process instances.

Process Model The process model is a pair $\mathcal{P} = \langle N, L \rangle$, comprising a directed multigraph N, the *process net*, and a *labeling function* L, defined on the edges of N.

The net $N = \langle P, T \rangle$ consists of a set of vertices P, the *places*, and a multiset of edges T, the *transitions*. A process instance traverses the net, starting from the distinguished place *start*. The transitions emanating from a place represent alternative developments of an instance.

A process instance has input data associated with it, which are represented by a fact $In(\bar{c}, \tau)$, where In is distinguished relation symbol, $\bar{c}$ is a tuple of constants from $dom_{\mathbb{Q}^+}$, and $\tau \in \mathbb{Q}^+$ is a time stamp that records when the process instance was started. We denote with $\Sigma_{\mathcal{B},In}$ and $\Sigma_{\mathcal{B}}$ the schemas of $\mathcal{B}$ with and without In, respectively.

The labeling function L assigns to every transition $t \in T$ a pair $L(t) = (E_t, W_t)$. Here, E_t, the *execution condition,* is a Boolean query over $\Sigma_{\mathcal{B},In}$ and W_t, the *writing rule,* is a rule $R(\bar{u}) \leftarrow B_t(\bar{u})$ whose head is a relation of $\Sigma_{\mathcal{B}}$ and whose body is a $\Sigma_{\mathcal{B},In}$-query that has the same arity as the head relation. Evaluating W_t over a $\Sigma_{\mathcal{B},In}$-instance $\mathcal{D}$ results in the set of facts $W_t(\mathcal{D}) = \{R(\bar{c}) \mid \bar{c} \in B_t(\mathcal{D})\}$. Intuitively, E_t specifies in which state of the database which process instance can perform the transition t, and W_t specifies which new information is (or can be) written into the database when performing t. In this paper we assume that E_t and B_t are conjunctive queries, possibly with negated atoms and inequality atoms with "$<$" and "$\leq$" involving timestamps. We assume inequalities to consist of one constant and one variable, like $X < 1^{\text{st}}$ Sep. We introduce these restricted inequalities so that we can model deadlines, without introducing an additional source of complexity for reasoning.

Configuration This component models the dynamics of an MDBP. Formally, a configuration is a triple $\langle \mathcal{I}, \mathcal{D}, \tau \rangle$, consisting of a part $\mathcal{I}$ that captures the process instances, a database instance $\mathcal{D}$ over $\Sigma_{\mathcal{B}}$, and a timestamp τ, the current time. The instance part, again, is a triple $\mathcal{I} = \langle O, M_{In}, M_P \rangle$, where $O = \{o_1, ..., o_k\}$ is a set of objects, called *process instances*,

and M_{In}, M_P are mappings, associating each $o \in O$ with a fact $M_{In}(o) = In(\bar{c}, \tau)$, its input record, and a place $M_P(o) \in P$, its current, respectively.

The input record is created when the instance starts and cannot be changed later on. While the data of the input record may be different from the constants in the database, they can be copied into the database by writing rules. A process instance can see the entire database, but only its own input record.

For convenience, we also use the notation $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D}, \tau \rangle$, $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ (when τ is not relevant), or $\mathcal{B} = \langle \mathcal{P}, \mathcal{D} \rangle$ (for a process that is initially without running instances).

Execution of an MDBP Let $\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$ be an MDBP, with current configuration $\mathcal{C} = \langle \mathcal{I}, \mathcal{D}, \tau \rangle$. There are two kinds of *atomic execution* steps of an MDBP: *(i)* the *traversal* of a transition by an instance and *(ii)* the *start* of a new instance.

***(i)* Traversal** of a transition. Consider an instance $o \in O$ with record $M_{In}(o) = In(\bar{c}, \tau')$, currently at place $M_P(o) = q$. Let t be a transition from q to p, with execution condition E_t. Then t is *enabled for* o, i.e., o can *traverse* t, if E_t evaluates to *true* over the database $\mathcal{D} \cup \{In(\bar{c}, \tau')\}$. Let $W_t \colon R(\bar{u}) \leftarrow B_t(\bar{u})$ be the writing rule of t. Then the effect of o traversing t is the transition from $\mathcal{C} = \langle \mathcal{I}, \mathcal{D}, \tau \rangle$ to a new configuration $\mathcal{C}' = \langle \mathcal{I}', \mathcal{D}', \tau \rangle$, such that *(i)* the set of instances O and the current time τ are the same; *(ii)* the new database instance is $\mathcal{D}' = \mathcal{D} \cup W_t(\mathcal{D} \cup \{In(\bar{c}, \tau')\})$, and *(iii)* $\mathcal{I} = \langle O, M_{In}, M_P \rangle$ is updated to $\mathcal{I}' = \langle O, M_{In}, M'_P \rangle$ reflecting the change of place for the instance o, that is $M'_P(o) = p$ and $M'_P(o') = M_P(o')$ for all other instances o'.

***(ii)* Start** of a new instance. Let o' be a fresh instance and let $In(\bar{c}', \tau')$ be an In-fact with $\tau' \geq \tau$, the current time of $\mathcal{C}$. The result of starting o' with info $\bar{c}'$ at time τ' is the configuration $\mathcal{C}' = \langle \mathcal{I}', \mathcal{D}, \tau' \rangle$ where $\mathcal{I}' = \langle O', M'_{In}, M'_P \rangle$ such that *(i)* the database instance is the same as in $\mathcal{C}$, *(ii)* the set of instances $O' = O \cup \{o'\}$ is augmented by o', and *(iii)* the mappings M'_{In} and M'_P are extensions of M_{In} and M_P, resp., obtained by defining $M'_{In}(o') = In(\bar{c}', \tau')$ and $M'_P(o') = start$.

An *execution* Υ of $\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$ is a finite sequence of configurations $\mathcal{C}_1, \ldots, \mathcal{C}_n$ *(i)* starting with $\mathcal{C}$ $(= \mathcal{C}_1)$, where *(ii)* each $\mathcal{C}_{i+1}$ is obtained from $\mathcal{C}_i$ by an atomic execution step. We denote Υ also with $\mathcal{C}_1 \rightsquigarrow \cdots \rightsquigarrow \mathcal{C}_n$. We say that the execution Υ *produces* the facts $A_1, \ldots, A_n$ if the database of the last configuration $\mathcal{C}_n$ in Υ contains $A_1, \ldots, A_n$. Since at each step a new instance can start, or an instance can write new data, *(i)* there are infinitely many possible executions, and *(ii)* the database may grow in an unbounded way over time.

3 The Query Stability Problem

In this section, we define the problem of query stability in MDBPs with its variants.

▶ **Definition 1** (Query Stability)**.** Given $\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$ with database instance $\mathcal{D}$, a query Q, and a timestamp τ, we say that Q *is stable in* $\mathcal{B}$ *until* τ, if for every execution $\mathcal{C} \rightsquigarrow \cdots \rightsquigarrow \mathcal{C}'$ in $\mathcal{B}$, where $\mathcal{C}'$ has database $\mathcal{D}'$ and timestamp τ' such that $\tau' < \tau$, it holds that

$$Q(\mathcal{D}) = Q(\mathcal{D}').$$

If the query Q is stable until time point ∞, we say it is *globally stable*, or simply, *stable*.

The interesting question from an application view is: *Given an MDBP $\mathcal{B}$, a query Q, and a timestamp τ, is Q stable in $\mathcal{B}$ until τ?* Stability until a time-point τ can be reduced to global stability. One can modify a given MDBP by adding a new *start* place and connecting it to the old *start* place via a transition that is enabled only for instances with timestamp

Defining Restrictions	Optional Restrictions	Data	Instance	Process	Query	Combined	Sect.
—	*fresh*†,*acyclic*	Undec.	Undec.	Undec.	Undec.	Undec.	4
closed	—	co-NP	co-NP	co-NExpTime	Π_2^p	co-NExpTime	6
positive	*closed*	PTime	co-NP	ExpTime	Π_2^p	ExpTime	5, 8
positive	*fresh*†, *acyclic*	PTime	co-NP	ExpTime	Π_2^p	ExpTime	8
closed, acyclic	*positive*	in AC^0	co-NP	PSpace	Π_2^p	PSpace	7
rowo	$*^\dagger$	in AC^0	in AC^0	co-NP	Π_2^p	Π_2^p	9

Table 1 Computational complexity of query stability in MDBPs. The results in a row hold for the class of MDBPs satisfying the defining restrictions and for the subclasses satisfying one or more of the optional restrictions. The results for all decidable variants indicate matching lower and upper bounds (except for AC^0). The $*$ indicates that the results for rowo hold for all non-trivial combinations of restrictions. All results for data, process, query and combined complexity of the decidable variants hold already for singleton MDBPs. †Note that, in all fresh variants instance complexity can be trivially decided in constant time (omitted in the table).

smaller than τ. Then a query Q is globally stable in the resulting MDBP iff in the original MDBP it is stable until τ.

To investigate sources of complexity and provide suitable encodings into Datalog, we identify five restrictions on MDBPs.

▶ **Definition 2** (Restriction on MDBPs and MDBP Executions)**.** Let $\mathcal{B}$ be an MDBP.

Positive: $\mathcal{B}$ is *positive* if execution conditions and writing rules contain only positive atoms;

Fresh: $\mathcal{B}$ is *fresh* if its configuration does not contain any running instances;

Acyclic: $\mathcal{B}$ is *acyclic* if the process net is cycle-free;

Rowo: $\mathcal{B}$ is *rowo* (= read-only-write-only) if the schema Σ of $\mathcal{B}$ can be split into two disjoint schemas: the reading schema Σ_r and the writing schema Σ_w, such that execution conditions and queries in the writing rules range over Σ_r while the heads range over Σ_w;

Closed: an execution of $\mathcal{B}$ is *closed* if it contains only transition traversals and no new instances are started.

We will develop methods for stability checking in MDBPs for all combinations of those five restrictions. For convenience, we will say that an MDBP $\mathcal{B}$ is *closed* if we consider only closed executions of $\mathcal{B}$. A *singleton MDBP* is a closed MDBP with a single instance in the initial configuration.

Complexity Measures The input for our decision problem are an MDBP $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$, consisting of a process model $\mathcal{P}$, an instance part $\mathcal{I}$, a database $\mathcal{D}$ and a timestamp τ, and a query Q. The question is: *Is Q globally stable in $\langle \mathcal{P}, \mathcal{I}, \mathcal{D}, \tau \rangle$?* We refer to process, instance, data, and query complexity if all parameters are fixed, except the process model, the instance part, the database, or the query, respectively.

Roadmap As a summary of our results, Table 1 presents the complexity of the possible variants of query stability. Each section of the sequel will cover one row.

Datalog Notation We assume familiarity with Datalog concepts such as *least fix point* and *stable model* semantics, and *query answering* over Datalog programs under both semantics. We consider Datalog programs that are *recursive*, *non-recursive*, *positive*, *semipositive*, *with negation*, or *with stratified negation* [9]. We write $\Pi \cup \mathcal{D}$ to denote a program where Π is a set of rules and $\mathcal{D}$ is a set of facts.

Notation	Meaning
$\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$, $\langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$, $\langle \mathcal{P}, \mathcal{I}, \mathcal{D}, \tau \rangle$	MDBP
$\mathcal{P} = \langle N, L \rangle$	Process part $\mathcal{P}$ with process net N and labeling function L
$\mathcal{C} = \langle \mathcal{I}, \mathcal{D}, \tau \rangle$	Configuration $\mathcal{C}$ with instance part $\mathcal{I}$, database instance $\mathcal{D}$ and timestamp τ
$N = \langle P, T \rangle$	Process net N with places P and transitions T
$T = \{t_1, \ldots, t_m\}$	Multi set of process transitions
p, q	Process places
$L(t) = (E_t, W_t)$	Labeling $L(t)$ of transition t
$In(\bar{s}, \tau), In(\bar{s})$	Input records with and without timestamps
$E_t \colon In(\bar{s}), S_1, \ldots, S_n, G_t$	Execution condition with possibly negated atoms S_i and conjunction of inequalities G_t
$W_t \colon R(\bar{u}) \leftarrow B_t(\bar{u})$	Writing rule with query $B_t(\bar{u}) \leftarrow In(\bar{s}), S_1, \ldots, S_l, M_t$ with possibly negated atoms S_i and conjunction of inequalities M_t
$\mathcal{I} = \langle O, M_{In}, M_P \rangle$	Instance part with process instances $O = \{o_1, \ldots, o_k\}$
$\Upsilon \colon \mathcal{C}_1 \rightsquigarrow \cdots \rightsquigarrow \mathcal{C}_n$	MDBP execution with configurations $\mathcal{C}_1, \ldots, \mathcal{C}_n$
r, a, k, m, c	MDBP $\mathcal{B}$ has r relations with a as the maximal arity, k running instances, m transitions, c number of constants

Table 2 Notation table of symbols that represent MDBPs

Summary of notation For convenience, we summarize the notation of our model used in the following sections in Table 2.

4 Undecidable MDBPs

With negation in execution conditions and writing rules, we can create MDBPs that simulate Turing machines (TMs). Consequently, in the general variant query stability is undecidable.

Due to lack of space we only provide an intuition. To show undecidability in data complexity, we define a database schema that allows us to store a TM and we construct a process model that simulates the executions of the stored TM. MDBPs cannot update facts in the database. However, we can augment relations with an additional version argument and simulate updates by adding new versions of facts. Exploiting negation in conditions and rules we can then refer to the last version of a fact. To simulate the TM execution, the process model uses fresh constants to model *(i)* an unbounded number of updates of the TM configurations (= number of execution steps in the TM), and *(ii)* a potentially infinite tape. The TM halts iff the process produces the predicate *dummy*. Undecidability in process complexity follows from undecidability in data complexity, since a process can first write the encoding of the TM into an initially empty database. Similarly, we obtain undecidability in instance complexity using instances that write the encoding of the TM at the beginning. To obtain undecidability in query complexity we extend the encoding for data complexity such that the database encodes a universal TM and an input of the TM is encoded in the query.

▶ **Theorem 3** (Undecidability). *Query stability in MDBPs is undecidable in data, process and query complexity. It is also undecidable in instance complexity except for fresh variants for which it is constant. Undecidability already holds for acyclic MDBPs.*

In our reduction it is the unbounded number of fresh instances that are causing writing rules to be executed an unbounded number of times, so that neither cycles nor existing instances are contributing to undecidability. In the sequel we study MDBPs that are positive, closed, or rowo, and show that in all three variants stability is decidable.

5 Positive Closed MDBPs

In cyclic positive MDBPs, executions can be arbitrarily long. Still, in the absence of fresh instances, it is enough to consider executions of bounded length to check stability. Consider a positive MDBP $\mathcal{B} = \langle \mathcal{P}, \mathcal{C} \rangle$, possibly with cycles and disallowing fresh instances to start, with c different constants, r relations, k running instances, m transitions and a as the maximal arity of a relation in $\mathcal{P}$. We observe: *(i)* For each relation R in $\mathcal{P}$ there are up to $c^{arity(R)}$ new R-facts that $\mathcal{B}$ can produce. Thus, $\mathcal{B}$ can produce up to rc^a new facts in total. *(ii)* It is sufficient to consider executions that produce at least one new fact each mk steps. An execution that produces no new facts in mk steps has at least one instance that in those mk steps visits the same place twice without producing a new fact; those steps can be canceled without affecting the facts that are produced. *(iii)* Hence, it is sufficient to consider executions of maximal length $mkrc^a$.

Among these finitely many executions, it is enough to consider those that produce a maximal set of new facts. Since a process instance may have the choice among several transitions, there may be several such maximal sets. We identify a class of executions in positive closed MDBPs, called *greedy executions*, that produce all maximal sets.

Greedy Executions Intuitively, in a greedy execution instances traverse all cycles in the net in all possible ways and produce all that can be produced before leaving the cycle. To formalize this idea we identify two kinds of execution steps: *safe steps* and *critical steps*. A safe step is an execution step of an instance after which, given the current state of the database, the instance can return to its original place. A critical step is an execution step that is not safe. Based on this, we define *greedy sequences* and *greedy executions*. A greedy sequence is a sequence of safe steps that produces the largest number of new facts possible. A greedy execution is an execution where greedy sequences and critical steps alternate.

Let Υ be a greedy execution with i alternations of greedy sequences and critical steps. In the following, we characterize which are the transitions that instances traverse in the $i+1$-th greedy sequence and then in the $i+1$-th critical step. For a process instance o and the database D_Υ produced after Υ we define the *enabled graph* $N_{\Upsilon,o}$ as the multigraph whose vertices are the places of N (i.e., the process net of $\mathcal{B}$) and edges those transitions of N that are enabled for o given database D_Υ. Let $SCC(N_{\Upsilon,o})$ denote the set of strongly connected components (SCCs) of $N_{\Upsilon,o}$. Note that two different instances may have different enabled graphs and thus different SCCs. For a place p, let $N^p_{\Upsilon,o}$ be the SCC in $SCC(N_{\Upsilon,o})$ that contains p. Suppose that o is at place p after Υ. Then in the next greedy sequence, each instance o traverses the component $N^p_{\Upsilon,o}$ in all possible ways until no new facts can be produced, meaning that all instances traverse in an arbitrary order. Conversely, the next critical step is an execution step where an instance o traverses a transition that is not part of $N^p_{\Upsilon,o}$, and thus it leaves the current SCC. We observe that when performing safe transitions new facts may be written and new transitions may become executable. This can make SCCs of $N_{\Upsilon,o}$ to grow and merge, enabling new safe steps. With slight abuse of notation we denote such maximally expanded SCCs with $N_{\Upsilon,o}$, and with $N^p_{\Upsilon,o}$ the maximal component that contains p.

Properties of Greedy Executions We identify three main properties of greedy executions.

- A greedy execution is characterized by its critical steps, because an instance may have to choose one among several possible critical steps. In contrast, how safe steps compose a greedy sequence is not important for stability because all greedy sequences produce the same (maximal) set of facts.
- A greedy execution in an MDBP with m transitions and k instances can have at most mk critical steps. The reason is that an execution step can be critical only the first time it is executed, and any time after that it will be a safe step.
- Each execution can be transformed into a greedy execution such that if a query is instable in the original version then it is instable also in the greedy version. In fact, an arbitrary execution has at most mk critical steps. One can construct a greedy version starting from those critical steps, such that the other steps are part of the greedy sequences.

▶ **Lemma 4.** *For each closed execution Υ in a positive MDBP $\mathcal{B}$ that produces the set of ground atoms W, there exists a greedy execution Υ' in $\mathcal{B}$ that also produces W.*

Therefore, to check stability it is enough to check stability over greedy executions. In the following we define Datalog rules that compute facts produced by greedy executions.

Encoding into Datalog Let $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ be a positive MDBP with m transitions and k instances. Since critical steps uniquely characterize a greedy execution, we use a tuple of size up to mk to encode them. For example, if in a greedy execution Υ at the first critical step instance o_{l_1} traverses transition t_{h_1}, in the second o_{l_2} traverses t_{h_2}, and so on up to step i, we encode this with the tuple

$$\bar{\omega} = \langle o_{l_1}, t_{h_1}, \ldots, o_{l_i}, t_{h_i} \rangle.$$

Next, we define the relations used in the encoding. *(i)* For each relation R in $\mathcal{P}$ we introduce relations R^i (for i up to mk) to store all R-facts produced by an execution with i critical steps. Let Υ be the execution from above and let $\langle o_{l_1}, t_{h_1}, \ldots, o_{l_i}, t_{h_i} \rangle$ be the tuple representing it. Then, a fact of relation R^i has the form $R^i(o_{l_1}, t_{h_1}, \ldots, o_{l_i}, t_{h_i}; \bar{s})$, and it holds iff Υ produces the fact $R(\bar{s})$. Later on we use $\bar{\omega}$ to represent the tuple $\langle o_{l_1}, t_{h_1}, \ldots \rangle$. Facts of R^i are then represented as $R^i(\bar{\omega}; \bar{s})$. For convenience, we use a semicolon (;) instead of a comma (,) to separate encodings of different types in the arguments. *(ii)* To record the positions of instances after each critical step we introduce relations $State^i$ such that $State^i(\bar{\omega}; p_1, \ldots, p_k)$ encodes that after Υ is executed, instance o_1 is at p_1, o_2 is at p_2, and so on. *(iii)* To store the SCCs of the enabled graph we introduce relations SCC^i such that for a process instance o and a place p, the transition t belongs to $N^p_{\Upsilon,o}$ iff $SCC^i(\bar{\omega}; o, p, t)$ is true. *(iv)* To compute the relations SCC^i, we first need to compute which places are reachable by an instance o from place p. For that we introduce auxiliary relations $Reach^i$ such that in the enabled graph $N_{\Upsilon,o}$ instance o can reach place p' from p iff $Reach^i(\bar{\omega}; o, p, p')$ is true. *(v)* Additionally, we introduce the auxiliary relation In^0 that associates instances with their In-records, that is $In^0(o; \bar{s})$ is true iff the instance o has input record $In(\bar{s})$. With slight abuse of notation, we use $\bar{\omega}$ to denote also the corresponding greedy closed execution Υ.

In the following we define a Datalog program that computes the predicates introduced above for all possible greedy executions. The program uses stratified negation.

Initialization For each relation R in $\mathcal{P}$ we introduce the *initialization rule* $R^0(\overline{X}) \leftarrow R(\overline{X})$ to store what holds before any critical step is made. Then we add the fact rule $State^0(p_1, \ldots, p_k) \leftarrow \mathit{true}$ if in the initial configuration o_1 is at place p_1, o_2 at p_2, and so on.

Greedy Sequence: Traversal Rules Next, we introduce rules that compute enabled graphs. The relation $Reach^i$ contains the transitive closure of the enabled graph $N_{\bar{\omega},o}$ for each o and $\bar{\omega}$, encoding a greedy execution of length i. First, a transition t from q to be p gives rise to an edge in the enabled graph $N_{\bar{\omega},o}$ if instance o can traverse that t:

$$Reach^i(\overline{W}; O, q, p) \leftarrow E_t^i(\overline{W}; O).$$

Here, $E_t^i(\overline{W}; O)$ is a shorthand for the condition obtained from E_t by replacing $In(\bar{s})$ with $In^0(O; \bar{s})$ and by replacing each atom $R(\bar{v})$ with $R^i(\overline{W}; \bar{v})$. The tuple $\overline{W}$ consists of $2i$ many distinct variables to match every critical execution with i steps. It ensures that only facts produced by $\overline{W}$ are considered. The transitive closure is computed with the following rule:

$$Reach^i(\overline{W}; O, P_1, P_3) \leftarrow Reach^i(\overline{W}; O, P_1, P_2), Reach^i(\overline{W}; O, P_2, P_3).$$

Based on $Reach^i$, SCC^i is computed by including every transition t from q to p that an instance can reach, traverse, and from where it can return to the current place:

$$SCC^i(\overline{W}; O, P, t) \leftarrow \ Reach^i(\overline{W}; O, P, q), E_t^i(\overline{W}; O), Reach^i(\overline{W}; O, p, P).$$

Critical Steps: Traversal Rules We now want to record how an instance makes a critical step. An instance o_j can traverse transition t from q to p at the critical step $i+1$ if *(i)* o_j is at some place in $N^q_{\bar{\omega},o_j}$ at step i, *(ii)* it satisfies the execution condition E_t, *(iii)* and by traversing t it leaves the current SCC. The following *traversal rule* captures this:

$$\begin{aligned} &State^{i+1}(\overline{W}, o_j, t; P_1, \ldots, P_{j-1}, p, P_{j+1}, \ldots, P_k) \leftarrow \\ &\quad State^i(\overline{W}; P_1, \ldots, P_{j-1}, P, P_{j+1}, \ldots, P_k), Reach^i(\overline{W}; o_j, P, q), Reach^i(\overline{W}; o_j, q, P), \quad (1) \\ &\quad E_t^i(\overline{W}; o_j), \neg SCC^i(\overline{W}; o_j, P, t). \quad (2) \end{aligned}$$

Here, the condition *(i)* is encoded in line (1), and *(ii)* and *(iii)* are encoded in line (2).

Generation Rules A fact in R^{i+1} may hold because *(i)* it has been produced by the current greedy sequence or by the last critical step, or *(ii)* by some of the previous sequences or steps. Facts produced by previous sequences or steps are propagated with the *copy rule*: $R^{i+1}(\overline{W}, O, T; \overline{X}) \leftarrow State^{i+1}(\overline{W}, O, T; _), R^i(\overline{W}; \overline{X})$, copying facts $R(\overline{X})$ holding after $\overline{W}$ to all extensions of $\overline{W}$.

Then we compute the facts produced by the next greedy sequence. For each instance o_j, being at some place p_j after the last critical step in $\bar{\omega}$, and for each transition t that is in $N^{p_j}_{\bar{\omega},o_j}$, with writing rule $R(\bar{u}) \leftarrow B_t(\bar{u})$, we introduce the following *greedy generation rule*:

$$R^i(\overline{W}; \bar{u}) \leftarrow \ State^i(\overline{W}; _, \ldots, _, P_j, _, \ldots, _), SCC^i(\overline{W}; o_j, P_j, t), B_t^i(\overline{W}; o_j; \bar{u}),$$

where condition $B_t^i(\overline{W}; O; \bar{u})$ is obtained similarly as $E_t^i(\overline{W}; O)$. In other words, all transitions t that are in $N^{p_j}_{\bar{\omega},o_j}$ are fired simultaneously, and this is done for all instances.

The facts produced at the next critical step by traversing t, which has the writing rule $R(\bar{u}) \leftarrow B_t(\bar{u})$, are generated with the *critical generation rule*: $R^{i+1}(\overline{W}, O, t; \bar{u}) \leftarrow State^{i+1}(\overline{W}, O, t; _), B_t^i(\overline{W}; O; \bar{u})$.

Let $\Pi^{po,cl}_{\mathcal{P},\mathcal{I}}$ be the program encoding the positive closed $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ as described above.

▶ **Lemma 5.** *Let $\bar{\omega}$ be a greedy execution in the positive closed $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ of length i and $R(\bar{s})$ be a fact. Then $R(\bar{s})$ is produced by $\bar{\omega}$ iff $\Pi^{po,cl}_{\mathcal{P},\mathcal{I}} \cup \mathcal{D} \models R^i(\bar{\omega}; \bar{s})$.*

Test Program Now we want to test the stability of $Q(\overline{X}) \leftarrow R_1(\bar{u}_1), \ldots, R_n(\bar{u}_n)$. We collect all potential Q-answers using the relation Q'. A new query answer may be produced by an execution of any size i up to mk. Thus, for each execution of a size i from 0 to mk we introduce the Q'-rule

$$Q'(\overline{X}) \leftarrow R_1^i(\overline{W}; \bar{u}_1), \ldots, R_n^i(\overline{W}; \bar{u}_n). \tag{3}$$

Then, if there is a new query answer, the *test rule* "$\mathit{Instable} \leftarrow Q'(\overline{X}), \neg Q(\overline{X})$" fires the fact *Instable*. Let $\Pi^{test}_{\mathcal{P},\mathcal{I},Q}$ be the test program that contains Q, the Q'-rules, and the test rule.

▶ **Theorem 6.** *Q is instable in the positive closed $\mathcal{B}$ iff $\Pi^{po,cl}_{\mathcal{P},\mathcal{I}} \cup \mathcal{D} \cup \Pi^{test}_{\mathcal{P},\mathcal{I},Q} \models \mathit{Instable}$.*

Data and Process Complexity Since $\Pi^{po,cl}_{\mathcal{P},\mathcal{I}} \cup \mathcal{D} \cup \Pi^{test}_{\mathcal{P},\mathcal{I},Q}$ is a Datalog program with stratified negation, for which reasoning is as complex as for positive Datalog, we obtain as upper bounds ExpTime for process and combined complexity, and PTime for data complexity [9]. We show that these are also lower bounds, even for singleton MDBPs. This reduction can also be adapted for acyclic fresh MDBPs, which we study in Section 8.

▶ **Lemma 7.** *Stability is* ExpTime*-hard in process and* PTime*-hard in data complexity for*
a) positive singleton MDBPs under closed executions, and
b) positive acyclic fresh MDBPs.

Proof Sketch. *a)* We encode query answering over a Datalog program $\Pi \cup \mathcal{D}$ into stability checking. Let A be a fact. We construct a positive singleton MDBP $\langle \mathcal{P}^{po,cl}_{\Pi,A}, \mathcal{I}_0, \mathcal{D} \rangle$, where there is a transition for each rule and the single process cycles to produce the least fixed point (LFP) of the program. In addition, the MDBP inserts the fact *dummy* if A is in the LFP. Then test query $Q_{test} \leftarrow \mathit{dummy}$ is stable in $\langle \mathcal{P}^{po,cl}_{\Pi,A}, \mathcal{I}_0, \mathcal{D} \rangle$ iff $\Pi \cup \mathcal{D} \not\models A$.

b) Analogous, letting fresh instances play the role of the cycling singleton instance. ◀

Instance Complexity Instance complexity turns out to be higher than data complexity. Already for acyclic positive closed MDBPs it is co-NP-hard because *(i)* process instances may non-deterministically choose a transition, which creates exponentially many combinations, even in the acyclic variant; and *(ii)* instances may interact by reading data written by other instances.

▶ **Lemma 8.** *There exist a positive acyclic process model $\mathcal{P}_0$, a database $\mathcal{D}_0$, and a test query Q_{test} with the following property: for every graph G one can construct an instance part $\mathcal{I}_G$ such that G is not 3-colorable iff Q_{test} is stable in $\langle \mathcal{P}_0, \mathcal{I}_G, \mathcal{D}_0 \rangle$ under closed executions.*

Clearly, Lemma 8 implies that checking stability for closed MDBPs is co-NP-hard in instance complexity. According to Theorem 11 (Section 6), instance complexity is co-NP for all closed MDBPs, which implies co-NP-completeness even for the acyclic variant.

Query Complexity To analyze query complexity we first show how difficult it is to check whether a query returns the same answer over a database and an extension of that database.

▶ **Lemma 9** (Answer Difference)**.** *For every two fixed databases $\mathcal{D} \subseteq \mathcal{D}'$, checking whether a given conjunctive query Q satisfies $Q(\mathcal{D}) = Q(\mathcal{D}')$ is in Π^{P}_2 in the query size. Conversely, there exist databases $\mathcal{D}_0 \subseteq \mathcal{D}'_0$ such that checking for a conjunctive query Q whether $Q(\mathcal{D}_0) = Q(\mathcal{D}'_0)$ is Π^{P}_2-hard in the query size.*

Proof Idea. The first claim holds since one can check $Q(\mathcal{D}) \subsetneq Q(\mathcal{D}')$ in NP using an NP oracle. We show the second by reducing the *3-coloring extension* problem for graphs [2]. ◀

Building upon Lemma 9, we can define an MDBP that starting from $\mathcal{D}_0$ produces $\mathcal{D}'_0$. In fact, for such an MDBP it is enough to consider the simplest variants of rowo.

▶ **Proposition 10.** *Checking stability is* Π^{P}_2*-hard for*

a) positive fresh acyclic rowo MDBPs, and
b) positive closed acyclic rowo singleton MDBPs.

Given $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$, there are finitely many maximal extensions $\mathcal{D}'$ of $\mathcal{D}$ that can be produced by $\mathcal{B}$. We can check stability of a query Q by finitely many checks whether $Q(\mathcal{D}) = Q(\mathcal{D}')$. Since each such check is in Π^{P}_2, according to Lemma 9, the entire check is in Π^{P}_2. Thus, stability is Π^{P}_2-complete in query complexity.

6 Closed MDBPs

In the presence of negation, inserting new facts may disable transitions. During an execution, a transition may switch many times between being enabled and disabled, and greedy executions could have exponentially many critical steps. An encoding along the ideas of the preceding section would lead to a program of exponential size. This would give us an upper bound of double exponential time for combined complexity. Instead, we establish a correspondence between stability and brave query answering for Datalog with (unstratified) negation under *stable model semantics* (SMS) [9]. Due to lack of space we only state the results.

▶ **Theorem 11.** *For every closed MDBP* $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ *and every query* Q *one can construct a Datalog program with negation* $\Pi^{cl}_{\mathcal{P}}$*, based on* $\mathcal{P}$*, a database* $\mathcal{D}_{\mathcal{I}}$*, based on* $\mathcal{D}$ *and* $\mathcal{I}$*, and a test program* Π^{test}_{Q}*, based on* Q*, such that the following holds:*

$$Q \text{ is instable in } \mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle \quad \text{iff} \quad \Pi^{cl}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{I}} \cup \Pi^{test}_{Q} \models_{brave} \textit{Instable}.$$

Proof Idea. For the same reason as in the positive variant, it is sufficient to consider executions of maximal length $mkrc^a$. Program $\Pi^{cl}_{\mathcal{P}}$ contains two parts: *(i)* a program that generates a linear order of size $mkrc^a$ (with parameters m, k, r, c, a defined as in Section 5), starting from an exponentially smaller order, that is used to enumerate execution steps, and *(ii)* a program that "guesses" an execution of size up to $mkrc^a$ by selecting for each execution step one instance and one transition, and that produces the facts that would be produced by the guessed execution. Then each execution corresponds to one stable model. The test program Π^{test}_{Q} checks if any of the guessed executions yields a new query answer. ◀

In Theorem 11, the process is encoded in the program rules while data and instances are encoded as facts. Since brave reasoning under SMS is NExpTime in program size and NP in data size [9], we have that process and combined complexity are in co-NExpTime, and data and instance complexity are in co-NP. From this and Lemma 8 it follows that instance complexity is co-NP-complete. To show that stability is co-NExpTime-complete in process and co-NP-complete in data complexity we encode brave reasoning into stability. Query complexity is Π^{P}_2-complete for the same reasons as in the positive variant.

▶ **Theorem 12.** *For every Datalog program* $\Pi \cup \mathcal{D}$*, possibly with negation, and fact* A*, one can construct a singleton MDBP* $\langle \mathcal{P}_{\Pi,A}, \mathcal{I}_0, \mathcal{D} \rangle$ *such that for the query* $Q_{test} \leftarrow$ *dummy we have:* $\Pi \cup D \models_{brave} A$ *iff* Q_{test} *is stable in* $\langle \mathcal{P}_{\Pi,A}, \mathcal{I}_0, \mathcal{D} \rangle$ *under closed executions.*

7 Acyclic Closed MDBPs

If a process net is cycle-free, all closed executions have finite length. More specifically, in an acyclic MDBP with m transitions and k running instances, the maximal length of an execution is mk. Based on this observation, we modify the encoding for the positive closed variant in Section 5 so that it can cope with negation and exploit the absence of cycles.

For an acyclic MDBP, there cannot exist any greedy steps, which would stay in a strongly connected component of the net. Therefore, we drop the encodings of greedy traversals and the greedy generation rules. We keep the rules for critical steps, but drop the atoms of relations Reach^i and SCC^i. Now, in contrast to the positive closed variant, we may have negation in the conditions E_t and B_t. However, the modified Datalog program is non-recursive, since each relation R^i and State^i is defined in terms of R^j's and State^j's where $j < i$.

Let $\Pi^{ac,cl}_{\mathcal{P},\mathcal{I}}$ be the program encoding an acyclic $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ as described above and let $\Pi^{test}_{\mathcal{P},\mathcal{I},Q}$ be the test program as in the cyclic variant.

▶ **Theorem 13.** *Q is instable in the closed acyclic $\mathcal{B}$ iff* $\Pi^{ac,cl}_{\mathcal{P},\mathcal{I}} \cup \mathcal{D} \cup \Pi^{test}_{\mathcal{P},\mathcal{I},Q} \models \mathit{Instable}$.

Complexity As upper bounds for combined and data complexity, the encoding gives us the analogous bounds for non-recursive Datalog$^\neg$ programs, that is, PSPACE in combined and AC^0 in data complexity [9]. Already in the positive variant, we inherit PSPACE-hardness of process complexity (and therefore also of combined complexity) from the program complexity of non-recursive Datalog. We obtain matching lower bounds by a reverse encoding.

▶ **Lemma 14.** *For every non-recursive Datalog program Π and every fact A, one can construct a singleton acyclic positive MDBP $\langle \mathcal{P}_{\Pi,A}, \mathcal{C}_0 \rangle$ such that for the query $Q_{test} \leftarrow$ dummy we have: $\Pi \not\models A$ iff Q_{test} is stable in $\langle \mathcal{P}_{\Pi,A}, \mathcal{C}_0 \rangle$ under closed executions.*

We observe that for closed executions, the cycles increase the complexity, and moreover, cause a split between variants with and without negation. Lemma 8 and Theorem 11 together imply that instance complexity is CO-NP-complete. Query complexity is Π^{P}_2-complete for the same reasons as in other closed variants.

8 Positive Fresh MDBPs

All decidable variants of MDBPs that we investigated until now were so because we allowed only closed executions. In this and the next section we show that decidability can also be guaranteed if conditions and rules are positive, or if relations are divided into read and write relations (rowo). We look first at the case where initially there are no running instances.

When fresh instances start, their input can bring an arbitrary number of new constants into the database. Thus, processes can produce arbitrarily many new facts. First we show how infinitely many executions of a positive or rowo MDBP can be faithfully abstracted to finitely many over a simplified process such that a query is stable over the original process iff it is stable over the simplified one. For such simplified positive MDBPs, we show how to encode stability checking into query answering in Datalog.

Abstraction Principle Let $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D}, \tau_{\mathcal{B}} \rangle$ be a positive or rowo MDBP and let Q be a query that we want to check for stability. Based on $\mathcal{B}$ and Q we construct an MDBP $\mathcal{B}' = \langle \mathcal{P}', \mathcal{I}, \mathcal{D}, \tau_{\mathcal{B}} \rangle$ that has the same impact on the stability of Q but uses at most linearly many fresh values from the domain.

Let *adom* be the active domain of $\mathcal{B}$ and Q, that is the set of all constants appearing in $\mathcal{B}$ and Q. Let $\tau_1, \ldots, \tau_n$ be all timestamps including $\tau_{\mathcal{B}}$ that appear in comparisons in $\mathcal{B}$ such that $\tau_i < \tau_{i+1}$. We introduce $n+1$ many fresh timestamps $\tau'_0, \ldots, \tau'_n \notin adom$ such that $\tau'_0 < \tau_1 < \tau'_1 < \cdots < \tau_n < \tau'_n$. If there are no comparisons in $\mathcal{B}$ we introduce one fresh timestamp τ'_0. Further, let a be a fresh value such that $a \notin adom$. Let $adom^* = adom \cup \{\tau'_0, \ldots, \tau'_n\} \cup \{a\}$ be the *extended active domain*.

Then, we introduce the *discretization* function $\delta_{\mathcal{B}} \colon dom_{\mathbb{Q}^+} \to dom_{\mathbb{Q}^+}$ that based on $adom^*$ "discretizes" $dom_{\mathbb{Q}^+}$ as follows: for each $\tau \in \mathbb{Q}^+$ *(i)* $\delta_{\mathcal{B}}(\tau) = \tau$ if $\tau = \tau_i$ for some i; *(ii)* $\delta_{\mathcal{B}}(\tau) = \tau'_i$ if $\tau_i < \tau < \tau_{i+1}$ for some i; *(iii)* $\delta_{\mathcal{B}}(\tau) = \tau'_0$ if $\tau < \tau_1$; *(iv)* and $\delta_{\mathcal{B}}(\tau) = \tau'_n$ if $\tau_n < \tau$; *(v)* and for $c \in dom$ if $c \in adom^*$ then $\delta_{\mathcal{B}}(c) = c$; otherwise $\delta_{\mathcal{B}}(c) = a$. If $\mathcal{B}$ has no comparisons then $\delta_{\mathcal{B}}(\tau) = \tau'_0$ for each τ. We extend $\delta_{\mathcal{B}}$ to all syntactic objects containing constants, including executions. Now, we define $\mathcal{P}'$ to be as $\mathcal{P}$, except that we add conditions on each outing transition from *start* such that only instances with values from $adom^*$ can traverse, and instances with the timestamps greater or equal than $\tau_{\mathcal{B}}$.

▶ **Proposition 15** (Abstraction). *Let $\Upsilon = \mathcal{C} \rightsquigarrow \mathcal{C}_1 \rightsquigarrow \cdots \rightsquigarrow \mathcal{C}_m$ be an execution in $\mathcal{B}$ that produces a set of facts W, and let $\Upsilon' = \delta_{\mathcal{B}}\Upsilon = \delta_{\mathcal{B}}\mathcal{C} \rightsquigarrow \delta_{\mathcal{B}}\mathcal{C}_1 \rightsquigarrow \cdots \rightsquigarrow \delta_{\mathcal{B}}\mathcal{C}_m$. Further, let Υ'' be an execution in $\mathcal{B}'$. Then the following holds:*
a) Υ' is an execution in $\mathcal{B}'$ that produces $\delta_{\mathcal{B}}W$;
b) $Q(\mathcal{D}) \neq Q(\mathcal{D} \cup W)$ iff $Q(\mathcal{D}) \neq Q(\mathcal{D} \cup \delta_{\mathcal{B}}W)$;
c) Υ'' is an execution in $\mathcal{B}$.

In other words, each execution in $\mathcal{B}$ can be $\delta_{\mathcal{B}}$-abstracted and it will be an execution in $\mathcal{B}'$, and more importantly, an execution in $\mathcal{B}$ produces a new query answer if and only if the $\delta_{\mathcal{B}}$-abstracted version produces a new query answer in $\mathcal{B}'$.

Encoding into Datalog Since $\mathcal{B}'$ allows only finitely many new values in fresh instances, there is a bound on the maximal extensions of $\mathcal{D}$ that can be produced. Moreover, since there is no bound on the number of fresh instances that can start, there is only a single maximal extension of $\mathcal{D}$, say $\mathcal{D}'$, that can result from $\mathcal{B}'$. We now define the program $\Pi^{po,fr}_{\mathcal{P},Q} \cup \mathcal{D}$ whose least fixpoint is exactly this $\mathcal{D}'$.

First, we introduce the relations that we use in the encoding. To record which fresh instances can reach a place p in $\mathcal{P}$, we introduce for each p a relation In_p with the same arity as *In*. That is, $In_p(\bar{s})$ evaluates to true in the program iff an instance with the input record $In(\bar{s})$ can reach p. As in the closed variant, we use a primed version R' for each relation R to store R-facts produced by the process.

Now we define the rules. Initially, all relevant fresh instances (those with constants from $adom^*$) sit at the *start* place. We encode this by the *introduction rule*: $In_{start}(X_1, \ldots, X_n) \leftarrow adom^*(X_1), \ldots, adom^*(X_n)$. Here, with slight abuse of notation, $adom^*$ represents a unary relation that we initially instantiate with the constants from $adom^*$. Also initially, we make a primed copy of each database fact, that is, for each relation R in $\mathcal{P}$ we define the *copy rule*: $R'(\overline{X}) \leftarrow R(\overline{X})$.

Then we encode instance traversals. For every transition t that goes from a place q to a place p, we introduce a *traversal rule* that mimics how instances having reached q move on to p, provided their input record satisfies the execution condition for t. Let $E_t = In(\bar{s}), R_1(\bar{s}_1), \ldots, R_l(\bar{s}_l), G_t$ be the execution condition for t, where G_t comprises the comparisons. We define the condition $E'_t(\bar{s})$ as $In_q(\bar{s}), R'_1(\bar{s}_1), \ldots, R'_l(\bar{s}_l), G_t$, obtained from E_t by renaming the *In*-atom and priming all database relations. Then, the *traversal rule* for t is: $In_p(\bar{s}) \leftarrow E'_t(\bar{s})$. Here, $E'_t(\bar{s})$ is defined over the primed signature since a disabled transition may become enabled as new facts are produced.

Which facts are produced by traversing t is captured by a *generation rule*. Let W_t: $R(\bar{u}) \leftarrow B_t(\bar{u})$ be the writing rule for t, with the query $B_t(\bar{u}) \leftarrow In(\bar{s}'), R_1(\bar{s}'_1), \ldots, R_n(\bar{s}'_n), M_t$, where M_t comprises the comparisons. Define $B'_t(\bar{s}', \bar{u}) \leftarrow In_q(\bar{s}'), R'_1(\bar{s}'_1), \ldots, R'_n(\bar{s}'_n), M_t$. The corresponding generation rule is $R'(\bar{u}) \leftarrow E'_t(\bar{s}), B'_t(\bar{s}', \bar{u}), \bar{s} = \bar{s}'$, which combines the constraints on the instance record from E_t and W_t.

Let $\Pi^{po,fr}_{\mathcal{P},Q}$ be the program defined above, encoding the positive fresh $\mathcal{B}'$ obtained from $\mathcal{B}$. The program is constructed in such a way that it computes exactly the atoms that are in the maximal extension $\mathcal{D}'$ of $\mathcal{D}$ produced by $\mathcal{B}'$. Let $R'(\bar{v})$ be a fact.

▶ **Lemma 16.** *There is an execution in the positive fresh $\mathcal{B}$ producing $R(\bar{v})$ iff $\Pi^{po,fr}_{\mathcal{P},Q} \cup \mathcal{D} \models R'(\bar{v})$.*

Let Π^{test}_{Q} be defined like $\Pi^{test}_{\mathcal{P},\mathcal{I},Q}$ in Section 5, except that there is only one rule for Q', obtained from (3) by replacing R^i_j with R'_j. Then Proposition 15 and Lemma 16 imply:

▶ **Theorem 17.** *Q is instable the positive fresh $\mathcal{B}$ iff $\Pi^{po,fr}_{\mathcal{P},Q} \cup \mathcal{D} \cup \Pi^{test}_{Q} \models Instable$.*

Complexity Since $\Pi^{po,fr}_{\mathcal{P},Q} \cup \mathcal{D} \cup \Pi^{test}_{Q}$ is a program with stratified negation, stability checking over positive fresh MDBPs is in ExpTime for process and combined complexity, and in PTime for data complexity [9]. From Lemma 7 we know that these are also lower bounds for the respective complexity measures. Query complexity is Π^{P}_2-complete as usual, and instance complexity is trivial for fresh processes.

Positive MDBPs To reason about arbitrary positive MDBPs, we can combine the encoding for the fresh variant ($\Pi^{po,fr}_{\mathcal{P},Q}$) from this section and the one for the closed variant from Section 5 ($\Pi^{po,cl}_{\mathcal{P},\mathcal{I}}$). The main idea is that to obtain maximal extensions, each greedy execution sequence is augmented by also flooding the process with fresh instances. The complexities for the full positive variant are inherited from the closed variant.

9 Read-Only-Write-Only MDBPs

In general MDBPs, processes can perform recursive inferences by writing into relations from which they have read. It turns out that if relations are divided into read-only and write-only, the complexity of stability reasoning drops significantly.

The main simplifications in this case are that *(i)* one traversal per instance and transition suffices, since no additional fact can be produced by a second traversal; *(ii)* instead of analyzing entire executions, it is enough to record which paths an individual process instance can take and which facts it produces, since instances cannot influence each other. As a consequence, the encoding program can be non-recursive and it is independent of the instances in the process configuration. A complication arises, however, since the maximal extensions of the original database $\mathcal{D}$ by the MDBP $\mathcal{B}$ are not explicitly represented by this approach. They consist of unions of maximal extensions by each instance and are encoded into the test query, which is part of the program.

▶ **Theorem 18.** *For every rowo MDBP $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ and query Q one can construct a nonrecursive Datalog program $\Pi^{ro}_{\mathcal{P},Q}$, based on $\mathcal{P}$ and Q, and a database instance $\mathcal{D}_{\mathcal{I}}$, based on $\mathcal{D}$ and $\mathcal{I}$, such that: Q is instable in $\mathcal{B}$ iff $\Pi^{ro}_{\mathcal{P},Q} \cup \mathcal{D}_{\mathcal{I}} \models Instable$.*

From the theorem it follows that data and instance complexities are in AC^0, except for instance complexity in fresh variants, for which it is constant.

Process, Query and Combined Complexity Since CQ evaluation can be encoded into an execution condition, this gives us co-NP-hardness of stability in process complexity. We also show that it is in co-NP. First we note that due to the absence of recursion, one can check in NP whether a set of atoms is produced by a process instance.

▶ **Proposition 19.** *Let $\mathcal{B}$ be a singleton rowo MDBP. One can decide in* NP*, whether for given facts $A_1, \ldots, A_m$, there is an execution in $\mathcal{B}$ that produces $A_1, \ldots, A_m$.*

Next, suppose that $\mathcal{I}$, $\mathcal{D}$ and $Q(\bar{v}) \leftarrow B_1, \ldots, B_m$ are a fixed instance part, database and query. Given a process model $\mathcal{P}$, we want to check that Q is instable in $\mathcal{B}_\mathcal{P} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$. Making use of the abstraction principle for fresh constants, we can guess in polynomial time an instantiation $B'_1, \ldots, B'_n$ of the body of Q that returns an answer not in $Q(\mathcal{D})$. Then we verify that $B'_1, \ldots, B'_n$ are produced by $\mathcal{B}_\mathcal{P}$. Such a verification is possible in NP according to Proposition 19. We guess a partition of the set of facts $B'_1, \ldots, B'_n$, guess one instance, possibly fresh, for each component set of the partition, and verify that the component set is produced by the instance. Since all verification steps were in NP, the whole check is in NP.

Query complexity is Π_2^P-complete for the same reasons as in the general variant, and one can show that this is also the upper-bound for the combined complexity.

10 Related Work

Traditional approaches for business process modeling focus on the set of activities to be performed and the flow of their execution. These approaches are known as *activity-centric*. A different perspective, mainly investigated in the context of databases, consists in identifying the set of data (entities) to be represented and describes processes in terms of their possible evolutions. These approaches are known as *data-centric*.

In the context of activity-centric processes, Petri Nets (PNs) have been used for the representation, validation and verification of formal properties, such as absence of deadlock, boundedness and reachability [26, 27]. In PNs and their variants, a token carries a limited amount of information, which can be represented by associating to the token a set of variables, like in colored PNs [18]. No database is considered in PNs.

Among data-centric approaches, *Transducers* [1, 25] were among the first formalisms ascribing a central role to the data and how they are manipulated. These have been extended to *data driven web systems* [11] to model the interaction of a user with a web site, which are then extended in [10]. These frameworks express insertion and deletion rules using FO formulas. The authors verify properties expressed as FO variants of LTL, CTL and CTL* temporal formulas. The verification of these formulas results to be undecidable in the general case. Decidability is obtained under certain restrictions on the input, yielding to ExpSpace complexity for checking LTL formulas and co-NExpTime and ExpSpace for CTL and CTL* resp., in the propositional case.

Data-Centric Dynamic Systems (DCDSs) [4] describe processes in terms of guarded FO rules that evolve the database. The authors study the verification of temporal properties expressed in variants of μ-calculus (that subsumes CTL*-FO). They identify several undecidable classes and isolate decidable variants by assuming a bound on the size of the database at each step or a bound on the number of constants at each run. In these cases verification is ExpTime-complete in data complexity.

Overall, both frameworks are more general than MDBPs, since deletions and updates of facts are also allowed. This is done by rebuilding the database after each execution step. Further, our stability problem can be encoded as FO-CTL formula. However, our decidability

results for positive MDBPs are not captured by the decidable fragments of those approaches. In addition, the authors of the work above investigate the borders of decidability, while we focus on a simpler problem and study the sources of complexity. Concerning the process representation, both approaches adopt a rule-based specification. This makes the control flow more difficult to grasp, in contrast to activity-centric approaches where the control flow has an explicit representation.

Artifact-centric approaches [17] use artifacts to model business relevant entities. In [6, 14, 15] the authors investigate the verification of properties of artifact-based processes such as reachability, temporal constraints, and the existence of dead-end paths. However, none of these approaches explicitly models an underlying database. Also, the authors focus on finding suitable restrictions to achieve decidability, without a fine-grained complexity analysis as in our case.

Approaches in [3] and [5], investigate the challenge of combining processes and data, however, focusing on the problem of data provenance and of querying the process structure.

In [12, 20] the authors study the problem of determining whether a query over views is independent from a set of updates over the database. The authors do not consider a database instance nor a process. Decidability in rowo MDBPs can be seen as a special case of those.

In summary, our approach to process modeling is closer to the activity-centric one but we model manipulation of data like in the data-centric approaches. Also, having process instances and MDBPs restrictions gives finer granularity compared to data-centric approaches.

11 Discussion and Conclusion

Discussion An interesting question is how complex stability becomes if MDBPs are not monotonic, i.e., if updates or deletions are allowed. In particular, for positive MDBPs we can show the following. In acyclic positive closed MDBPs updates and deletions can be modeled using negation in the rules, thus stability stays PSpace-complete. For the cyclic positive closed variant, allowing updates or deletions is more powerful than allowing negation, and stability jumps to ExpSpace-completeness. For positive MDBPs with updates or deletions stability is undecidable.

In case the initial database is not known, our techniques can be still applied since an arbitrary database can be produced by fresh instances starting from an empty database.

Contributions Reasoning about data and processes can be relevant in decision support to understand how processes affect query answers. (1) To model processes that manipulate data we adopt an explicit representation of the control flow as in standard BP languages (e.g., BPMN). We specify how data is manipulated as annotations on top of the control flow. (2) Our reasoning on stability can be offered as a reasoning service on top of the query answering that reports on the reliability of an answer. Ideally, reasoning on stability should not bring a significant overhead on query answering in practical scenarios. Existing work on processes and data [4] shows that verification of general temporal properties is typically intractable already measured in the size of the data. (3) In order to identify tractable cases and sources of complexity we investigated different variants of our problem, by considering negation in conditions, cyclic executions, read access to written data, presence of pending process instances, and the possibility to start fresh process instances. (4) Our aim is to deploy reasoning on stability to existing query answering platforms such as SQL and ASP [19]. For this reason we established different encodings into suitable variants of Datalog, that are needed to capture the different characteristics of the problem. For each of them we

showed that our encoding is optimal. In contrast to existing approaches, which rely on model checking to verify properties, in our work we rely on established database query languages.

Open Questions In our present framework we cannot yet model process instances with activities that are running in parallel. Currently, we are able to deal with it only in case instances do not interact (like in rowo). Also, we do not know yet how to reason about expressive queries, such as conjunctive queries with negated atoms, and first-order queries. From an application point of view, stability of aggregate queries and aggregates in the process rules are relevant. A further question is how to quantify instability, that is, in case a query is not stable, how to compute the minimal and maximal number of possible new answers.

Acknowledgements This work was partially supported by the research projects MAGIC, funded by the province of Bozen-Bolzano, and CANDy and PARCIS, funded by the Free University of Bozen-Bolzano.

References

1 S. Abiteboul, V. Vianu, B.S. Fordham, and Y. Yesha. Relational Transducers for Electronic Commerce. In *PODS*, pages 179–187, 1998.

2 M. Ajtai, R. Fagin, and L.J. Stockmeyer. The Closure of Monadic NP. *J. Comput. Syst. Sci.*, 60(3):660–716, 2000.

3 Y. Amsterdamer, S. B. Davidson, D. Deutch, T. Milo, J. Stoyanovich, and V. Tannen. Putting Lipstick on Pig: Enabling Database-Style Workflow Provenance. *PVLDB*, 5(4):346–357, 2011.

4 B. Bagheri Hariri, D. Calvanese, G. De Giacomo, A. Deutsch, and M. Montali. Verification of Relational Data-Centric Dynamic Systems with External Services. In *PODS*, pages 163–174, 2013.

5 C. Beeri, A. Eyal, S. Kamenkovich, and T. Milo. Querying Business Processes. In *VLDB*, pages 343–354, 2006.

6 K. Bhattacharya, C. E. Gerede, R. Hull, R. Liu, and J. Su. Towards Formal Analysis of Artifact-Centric Business Process Models. In *BPM*, pages 288–304, 2007.

7 Bonitasoft. Bonita BPM. `www.bonitasoft.com`. Accessed: 2015-12-16.

8 G. Cong, W. Fan, F. Geerts, X. Jia, and S. Ma. Improving Data Quality: Consistency and Accuracy. In *VLDB*, pages 315–326, 2007.

9 E. Dantsin, T. Eiter, G. Gottlob, and A. Voronkov. Complexity and Expressive Power of Logic Programming. *ACM Comput. Surv.*, 33(3):374–425, 2001.

10 A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic Verification of Data-Centric Business Processes. In *ICDT*, pages 252–267, 2009.

11 A. Deutsch, L. Sui, and V. Vianu. Specification and Verification of Data-Driven Web Services. In *PODS*, pages 71–82, 2004.

12 C. Elkan. Independence of Logic Database Queries and Updates. In *PODS*, pages 154–160, 1990.

13 W. Fan, F. Geerts, and J. Wijsen. Determining the Currency of Data. *ACM Trans. Database Syst.*, 37(4):25, 2012.

14 C. E. Gerede, K. Bhattacharya, and J. Su. Static Analysis of Business Artifact-Centric Operational Models. In *SOCA*, pages 133–140, 2007.

15 C. E. Gerede and J. Jianwen Su. Specification and Verification of Artifact Behaviors in Business Process Models. In *ICSOC*, pages 181–192, 2007.

16 F. T. Heath, D. Boaz, M. Gupta, R. Vaculín, Y. Sun, R. Hull, and L. Limonad. Barcelona: A Design and Runtime Environment for Declarative Artifact-Centric BPM. In *ICSOC*, pages 705–709, 2013.

17 R. Hull. Artifact-Centric Business Process Models: Brief Survey of Research Results and Challenges. In *OTM*, pages 1152–1163, 2008.

18 K. Jensen and L.M. Kristensen. *Coloured Petri Nets: Modelling and Validation of Concurrent Systems*. Springer, 2009.

19 N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, and F. Scarcello. The DLV System for Knowledge Representation and Reasoning. *ACM Trans. Comput. Log.*, 7(3):499–562, 2006.

20 A.Y. Levy and Y. Sagiv. Queries Independent of Updates. In *VLDB*, pages 171–181, 1993.

21 E. Marengo, W. Nutt, and O. Savković. Towards a Theory of Query Stability in Business Processes. In *AMW*, volume 1189 of *CEUR Workshop Proceedings*, 2014.

22 Object Management Group. *Business Process Model and Notation 2.0 (BPMN)*, Jan 2011.

23 S. Razniewski and W. Nutt. Completeness of Queries over Incomplete Databases. *PVLDB*, 4(11):749–760, 2011.

24 O. Savković, E. Marengo, and W. Nutt. Query Stability in Data-aware Business Processes. Technical Report KRDB15-1, KRDB Research Center, Free Univ. Bozen-Bolzano, 2015. `http://www.inf.unibz.it/krdb/pub/tech-rep.php`.

25 M. Spielmann. Verification of Relational Transducers for Electronic Commerce. In *PODS*, pages 92–103. ACM, 2000.

26 W.M.P. van der Aalst. Verification of Workflow Nets. In *ICATPN*, pages 407–426, 1997.

27 W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow Patterns. *Distributed and Parallel Databases*, 14(1):5–51, 2003.

Part

Appendices

Table of Contents

A Example **20**
A.1 Scenario: Student Registration 20
A.2 MDBP Representation of the Scenario 20
A.3 Stability of the Example Queries 21
A.4 Variants of the Example Process 23

B Closed MDBPs **23**
B.1 Proof of Theorem 11 23
B.2 Proof of Proposition 12 25

C Rowo MDBPs **33**
C.1 Proof of Theorem 18 33
C.2 Proof of Proposition 19 37

A Example

As an illustration of the concepts in our formalism, we provide an example about student enrollment at a university.

A.1 Scenario: Student Registration

One year in November, the student office distributes a report on the numbers of new student registrations for the offered programs. When comparing the numbers with those of the previous years, the Master in Computer Science (*mscCS*) shows a decrease, in contrast with the Master in Economics (*mscEco*), which has registered a substantial increase. An analysis task force at university level cannot identify a plausible cause. Eventually, a secretary discovers that the reason is a complication in the registration process, which foresees two routes to registration: a regular one and a second one via international federated study programs to which some programs, like the *mscCS*, are affiliated. Due to different deadlines, regular registration has been concluded in November while registration for students from federated programs has not. Since the *mscCS* is affiliated to some federated programs, but the *mscEco* is not, the query asking for all *mscEco* students was stable in November, while the query for all *mscCS* students was not and returned too low a number.

A.2 MDBP Representation of the Scenario

Table 3 shows the student registration process $\mathcal{B}_{reg} = \langle \mathcal{P}_{reg}, \mathcal{C}_{reg} \rangle$. Part *(a)* reports the process net and Part *(b)* the execution conditions and writing rules.

A process instance starts when a student submits an online request, providing as input the student name S and the course name C. Automatically, the system attaches a time stamp $\mathcal{T}$ to the request. The application is then represented as an *In*-record $In(S, C, \mathcal{T})$.

The procedure distinguishes between applications to *international* courses, which are part of programs involving universities from different countries and where an international commission decides whom to admit, and to *regular* programs, where the university itself evaluates the applications. Accordingly, a process first checks for the type of application. The transition 'is intl. app.' can only be traversed, if the execution condition $\mathit{In}(S, C, \mathcal{T})$, $\mathit{StudyPlan}(C, \mathit{intl}, P)$ succeeds, which is the case when the course of the application is stored in the relation *StudyPlan* and associated to a program with type *intl*. Subsequently, the process checks if the student has already been admitted ('is admitted'). If so, it pursues the upper branch of the net. If not, it checks if the course is also open to regular applications ('isn't admitted'). Similarly, the execution condition on 'is reg. app.' ensures that the course is associated to a program of type *reg*, but not of type *intl*. Then, applications for regular courses follow the bottom branch.

Deadlines give rise to conditions on the application timestamp $\mathcal{T}$. While applications are accepted starting from 1st Sep, the deadline for regular courses is 31st Oct, and for international courses it is 31st Dec. Candidates who applied until 30th Sep can pre-enroll, that is, register provisionally. After that date, admitted candidates have to register directly.

Provisional registration gives students the possibility *(i)* to enroll conditionally and complete an application not fully complete, and *(ii)* to confirm or withdraw the registration before being formally enrolled. Modeling the completion of incomplete applications leads to cycles in the net, while non-determinism, e.g. due to human intervention or interaction with other systems, is modeled by labeling the transitions emanating from a place (like 'acad. check' or 'stud. decis.') with non-exclusive execution conditions.

Some transitions are labeled with a writing rule. When 'pre-enrol cond.' is traversed, the rule $\mathit{Conditional}(S, C) \leftarrow \mathit{In}(S, C, \mathcal{T})$ records that the application is conditionally accepted by writing a fact into the relation *Conditional*. This relation, on the other hand, is read by the execution condition of the transition 'complete app.'

Table 3*(c)* shows a database instance $\mathcal{D}_{reg}$ for our running example. Courses offered are stored in the relation *StudyPlan*, together with their type (*intl* or *reg*) and the program they are associated with. The remaining tables store information about the students. Table 3*(d)* reports the running process instances $\mathcal{I}_{reg}$ in the form of a relation.

A.3 Stability of the Example Queries

Consider the queries

$$Q_{\text{eco}}(S) \leftarrow \mathit{Registered}(S, C), \mathit{StudyPlan}(C, T, \mathit{mscEco}),$$
$$Q_{\text{cs}}(S) \leftarrow \mathit{Registered}(S, C), \mathit{StudyPlan}(C, T, \mathit{mscCS}),$$

which ask for the students registered for the master in CS, and the master in Economics, respectively. We analyze their stability over different periods, specified in Table 4. For each period from τ_1 to τ_2, we ask if the query is stable until τ_2 in a variant of $\mathcal{B}_{reg}$ where *(i)* the current date is in the interval and *(ii)* there are no running applications with a start date later than τ_2 (and also no data in the tables about the students having submitted one).

During the period before 1st Sep, neither program allows registrations to proceed and thus both queries are stable until this date. For the period 1st Sep–31st Nov, the programs allow for new registrations and both queries are instable. If the current time is within the period 1st Nov–31st Dec and there are no pending applications, Q_{eco} is stable because the program *mscEco* is not affiliated with any international (*intl*) course and the deadline for the regular programs has passed. However, *mscCS* has an affiliated course to which student

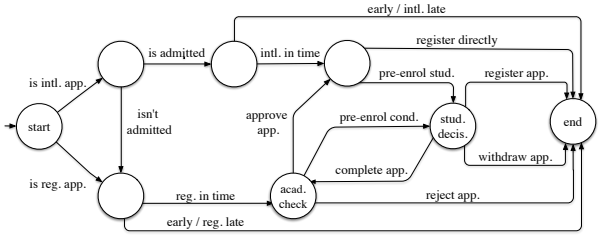

(a) Process Net in $\mathcal{P}_{reg}$

Transition	*Execution Condition (E)*
is intl. app.	$In(S, C, \mathcal{T})$, $StudyPlan(C, intl, P)$
is admitted	$In(S, C, \mathcal{T})$, $AdmittedIntl(S, C)$
isn't admitted	$In(S, C, \mathcal{T})$, $\neg AdmittedIntl(S, C)$, $StudyPlan(C, reg, P)$
is reg. app.	$In(S, C, \mathcal{T})$, $StudyPlan(C, reg, P)$, $\neg StudyPlan(C, intl, P)$
reg. in time	$In(S, C, \mathcal{T})$, 1^{st} Sep $\leq \mathcal{T} \leq 31^{st}$ Oct
intl. in time	$In(S, C, \mathcal{T})$, 1^{st} Sep $\leq \mathcal{T} \leq 31^{st}$ Dec
early	$In(S, C, \mathcal{T})$, $\mathcal{T} < 1^{st}$ Sep
reg. late	$In(S, C, \mathcal{T})$, $\mathcal{T} > 31^{st}$ Oct
intl. late	$In(S, C, \mathcal{T})$, $\mathcal{T} > 31^{st}$ Dec
pre-enrol stud.	$In(S, C, \mathcal{T})$, $\mathcal{T} \leq 30^{th}$ Sep
pre-enrol cond.	$In(S, C, \mathcal{T})$, $\mathcal{T} \leq 30^{th}$ Sep
register directly	$In(S, C, \mathcal{T})$, $\mathcal{T} > 30^{th}$ Sep
complete app.	$In(S, C, \mathcal{T})$, $Conditional(S, C)$
register app.	$In(S, C, \mathcal{T})$, $Pre\text{-}enrolled(S, C)$
withdraw app. = approve app. = reject app.:	*true*

(b) Execution Conditions and Writing Rules in $\mathcal{P}_{reg}$

StudyPlan

course	*type*	*program*
compLogic	intl	mscCS
compLogic	reg	mscCS
db	reg	mscCS
econ	reg	mscEco

AdmittedIntl

student	*course*
bob	compLogic
mary	compLogic

Pre-enrolled

student	*course*
bob	compLogic
alice	econ

Conditional

student	*course*
paul	econ

Registered

student	*course*
bob	compLogic
alice	econ

(c) Database Instance $\mathcal{D}_{reg}$

Instances and Mappings

id	*In*-record	place
o_1	(bob, compLogic, 5^{th} Sep)	*end*
o_2	(alice, econ, 20^{th} Sep)	*end*
o_3	(paul, econ, 24^{th} Sep)	*end*
o_4	(john, db, 4^{th} Nov)	*start*

(d) Process Instances $\mathcal{I}_{reg}$

Transition	*Writing Rule (W)*
register directly	$Registered(S, C) \leftarrow In(S, C, \mathcal{T})$
pre-enrol stud.	$Pre\text{-}enrolled(S, C) \leftarrow In(S, C, \mathcal{T})$
pre-enrol cond.	$Conditional(S, C) \leftarrow In(S, C, \mathcal{T})$
register app.	$Registered(S, C) \leftarrow In(S, C, \mathcal{T})$

(b) Execution Conditions and Writing Rules in $\mathcal{P}_{reg}$

Table 3 MDBP Representation of the Student Registration Process.

Period	$< 1^{st}$ Sep	1^{st} Sep–31^{st} Oct	1^{st} Nov–31^{st} Dec	$> 31^{st}$ Dec
Q_{eco}	stable	instable	stable	stable
Q_{cs}	stable	instable	instable	stable

Table 4 Stability of the queries Q_{eco} and Q_{cs} for different intervals.

Mary is admitted. She is not registered yet and potentially could submit an application before the 31^{st} Dec, which would be accepted. Thus, Q_{cs} is not stable for this period. If all the admitted students had already been registered, the query would be stable, since no new registration would be possible. The query would also be stable in the case the process is closed for new instances to start (e.g., because the limit on registered students has been reached). In this case, only running instances would be allowed to finish their execution. Thus, candidate *Mary* would not be able to register even though she is admitted. If the current time is after 31^{st} Dec, both queries are stable regardless whether the process is closed or not because all the registration deadlines have expired.

A.4 Variants of the Example Process

The model of $\mathcal{B}_{reg}$ is general, since the relations *Pre-enrolled* and *Conditional* are both read and written; the rules are normal, though only with negation on database relations that are not updated; the net is clearly cyclic. We can imagine that at the beginning of the registration period the process starts with a fresh configuration (i.e., no running applications). The case of arbitrary configurations includes situations that arise as exceptions in the registration process and cannot evolve from a fresh configuration. For instance, a regular application received after the deadline for a valid reason may be placed by a secretary at a certain place in the process that it would not be able to reach from the *start* place. Our example is not closed. If after the last deadline (31$^{\text{st}}$ Dec.) the web form for submitting new applications will be no more available the process will run under closed executions.

Note that, in our running example, negation in the conditions appears only on the database relations that are not updated by the process. For this case we can still apply the encoding from this section and obtain a semipositive Datalog program.

B Closed MDBPs

B.1 Proof of Theorem 11

The encoding program consists of the following rules: (a) *Ordering rules* that generate a linear order of size $mkrc^a$ (see Section 5) that we use to enumerate all executions steps; (b) *Selection rules* that for each execution step non-deterministically selects one instance and transition meaning that the selected instances traverses the selected transition at that step; (c) *Control rules* that discard cases where guesses execution sequence do not correspond to a valid execution in the process; (d) *Generation rules* that generate facts produced by a valid execution; (e) *Testing rules* that test if any of the guessed executions yield a new query answer.

Generating Exponentially Big Linear Order Assume we are given an MDBP $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ possible with cycles and negation in the rules. As we discussed, to check stability in $\mathcal{B}$ cyclic MDBPs it is sufficient to consider executions that have up to $mkrc^a$ executions steps, where m, k, r, c and a are parameters of $\mathcal{B}$ as defined in Section 5.

We introduce a Datalog program that generates a liner order of size $mkrc^a$ starting from a much smaller order (exponentially smaller). To define a small order we introduce a set of constants, called *digits*, $Dig_{\mathcal{B}} = \{d_1, \ldots, d_l\}$ of size l an we establish one linear order $<$ on $Dig_{\mathcal{B}}$: $d_1 < d_2 < \cdots < d_l$. Assume that $Dig^g_{\mathcal{B}}$ is the Cartesian power of $Dig_{\mathcal{B}}$ of size g. That is, each tuple $\bar{k}$ from $Dig^g_{\mathcal{B}}$ is of the form $\bar{k} = \langle d_{i_1}, \ldots, d_{i_g} \rangle$, for $d_{i_1}, \ldots, d_{i_g} \in Dig_{\mathcal{B}}$. We define $<^g$ as the lexicographical order on the tuples from $Dig^g_{\mathcal{B}}$. Here l and g are selected such that *(i)* $l = kc$ and thus depends on $\mathcal{P}, \mathcal{I}$ and D, and *(ii)* $g > m + r + a$ where m, r, and a are the parameters that depend only on $\mathcal{P}$. Then, it is not hard to check that it holds

$$l^g \geq mkrc^a.$$

In other words, linear order on $<^g$ is sufficient to enumerate all executions steps.

Adopting the idea in [9], we define a positive Datalog program that generates $<^g$. In particular, we want generate a relation *Succ* that stores the immediate successor in the order. The order $<^g$ is generated based on the orders $<^i$ on $Dig^i_{\mathcal{B}}$ for $i < g$. For that and to count the execution steps we introduce the the following relations.

Notation We introduce: *(i) Digit* – a unary relation such that $Digit(d)$ is true iff $d \in Dig_{\mathcal{B}}$; *(i) First*^{i} – an auxiliary relation of arity i such that $First^i(\bar{d})$ is true iff $\bar{d}$ is the first element

of the linear order $<^i$; *(ii)* $Last^i$ – an auxiliary relation of arity i such that $Last^i(\bar{d})$ is true iff $\bar{d}$ is the last element of the linear order $<^i$; *(ii)* $Step$ – a g-ary relation such that $Step(\bar{k})$ is true iff $\bar{k}$ is a tuple from $Dig^g_{\mathcal{B}}$ that corresponds to some execution step, that is any tuple from $Dig^g_{\mathcal{B}}$ except for the first one in the order $<^g$; *(iii)* $Succ^i$ – a relation of arity $2i$ such that $Succ^i(\bar{d}_1, \bar{d}_2)$ is true iff $\bar{d}_2$ is the immediate successor of $\bar{d}_1$ in the order $<^i$.

Ordering Rules To generate $Succ^i$, $First^i$, and $Last^i$ we introduce the following rules:

$$
\begin{aligned}
&Succ^{i+1}(Z, \overline{X}; Z, \overline{Y}) \leftarrow Digit(Z), Succ^i(\overline{X}; \overline{Y})\\
&Succ^{i+1}(Z_1, \overline{X}; Z_2, \overline{Y}) \leftarrow Succ^1(Z_1; Z_2), First^i(\overline{X}), Last^i(\overline{Y})\\
&First^{i+1}(X_1, \overline{X}) \leftarrow First^1(X_1), First^i(\overline{X})\\
&Last^{i+1}(Y_1, \overline{Y}) \leftarrow Last^1(Y_1), Last^i(\overline{Y}).
\end{aligned}
$$

Then, we populate relation $Step$ with the following rule:

$$Step(K_1, \ldots, K_g) \leftarrow Digit(K_1), \ldots, Digit(K_g), \neg First(K_1, \ldots, K_g)$$

In the following we use $Succ$ for $Succ^g$, $First$ for $First^g$, and $Last$ for $Last^g$.

We denote the above program as $\Pi^{succ}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{B}}$. Here, $\Pi^{succ}_{\mathcal{P}}$ is a program that is polynomial in the size of $\mathcal{P}$, and $\mathcal{D}_{\mathcal{B}}$ is a database instance that contains facts for relations $Digit$, $First^1$, $Last^1$ and $Succ^1$, and thus it is that is polynomial in the size of $\mathcal{B}$. Then it holds:

▶ **Lemma 20.** *Let $\bar{k}$, $\bar{k}_1$ and $\bar{k}_2$ be tuples of size g, then:*

(i) $\Pi^{succ}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{B}} \models Succ(\bar{k}_1, \bar{k}_2)$ *iff* $\bar{k}_2$ *is the successor of* $\bar{k}_1$.

(i) $\Pi^{succ}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{B}} \models Step(\bar{k})$ *iff* $\bar{k} \in Dig^g$;

In other words, the program $\Pi^{succ}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{B}}$ generates the linear order in PTime in the size of data and instances, and in ExpTime in the size of process.

Encoding Stability into Datalog with Negation In the following we define a Datalog program with negation that, based on the linear order from above, produces all maximal extended databases. Each maximal extended database is going to be encoded as one of the SMs of the program. The program adapts *guess and check* methodology from answer-set programming that organizes rules in guessing rules that generate SM candidates, and checking rules that discards bad candidates.

Notation To encode guessing of executions we introduce relations $Moved$ and $NotMoved$ of size $g+1$ such that $Moved(\bar{k}, o)$ means that instance o traverses at step $\bar{k}$, and $NotMoved(\bar{k}, o)$ means the opposite. Here, $NotMoved$ is needed for technical reasons. Similarly, for transitions we introduce relations $Trans$ and $NotTrans$ such that $Trans(\bar{k}, t)$ means that at step $\bar{k}$ transition t is traversed; and $NotTrans(\bar{k}, t)$ means the opposite. Further, we introduce relation $Completed$ of size g that we use to keep track of the steps that are completed. That is, $Completed(\bar{k})$ is true if step $\bar{k}$ is completed and steps that precede $\bar{k}$ are also completed. Then, to store the positions of each instance after some execution step we use relation $Place$. E.g., if after $\bar{k}$-th step instance o is at place p then $Place(\bar{k}, o, p)$ is true. To store facts that are produced up to a certain step we introduce prime version relation R' for each R in $\Sigma_{\mathcal{B}}$. Then $R'(\bar{k}, \bar{s})$ is true iff $R(\bar{s})$ is produced up to step $\bar{k}$.

First we define *guessing rules*:

$$
\begin{aligned}
Moved(\overline{K}; O) &\leftarrow Step(\overline{K}), In^0(O, _), \neg NotMoved(\overline{K}; O),\\
NotMoved(\overline{K}; O) &\leftarrow Step(\overline{K}), In^0(O, _), \neg Moved(\overline{K}; O),\\
\bot &\leftarrow Step(\overline{K}), Moved(\overline{K}; O_1), Moved(\overline{K}; O_2), O_1 \neq O_2,\\
\bot &\leftarrow In^0(O, _), Step(\overline{K}), \neg Moved(\overline{K}, O).
\end{aligned}
$$

Intuitively, the first two rules enforce each SM to partition instances into *Moved* and *NotMoved* for each step $\bar{k}$, and the last two ensures that at most one and at least one instance is selected.

We define the same kind of rules for *Trans* and *NotTrans*.

Once an instance and a transition have been selected for one execution step $\bar{k}$, we need to ensure that the instance can actually traverse the transition. Relation *Completed* keeps track of that for each step $\bar{k}$ by checking if *(i)* the selected instance o satisfies the execution condition of the selected transition t; *(ii)* the instance o is at place q from which t originates; and *(iii)* if all previous execution steps were already completed. This is achieved using the *checking rules*:

$$\begin{aligned} \mathit{Completed}(\overline{K}_2) \leftarrow\ & \mathit{Moved}(\overline{K}_2; O), \mathit{Trans}(\overline{K}_2; t), \mathit{Succ}(\overline{K}_1, \overline{K}_2), \\ & \mathit{Completed}(\overline{K}_1), E_t(\overline{K}_1; O), \mathit{Place}(\overline{K}_1; O; q). \end{aligned}$$

Condition $E_t(\overline{K}_1; O)$ is similar to the positive acyclic case where the execution $\bar{\omega}$ is replaced with the execution step $\overline{K}_1$.

Similarly, we define generation rules that for R' and rules that update position of instances store in *Place*. Let the above rules together with the program $\Pi^{succ}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{B}}$ define the program $\Pi^{cl}_{\mathcal{P}}$ for closed MDBPs, and let $\mathcal{D}_{\mathcal{I}}$ be a database instance that contains In^0 facts and $\mathcal{D}$.

▶ **Lemma 21.** *Let $\bar{k}$ be an execution step in $\mathcal{B}$, and let $R(\bar{s})$ be a fact. The following is equivalent:*

- *There is an execution of length $\bar{k}$ in $\mathcal{B}$ that produces $R(\bar{s})$;*
- $\Pi^{cl}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{I}} \models_{brave} R'(\bar{k}; \bar{s})$.

Now we want to test query Q for stability. We collect new query answers with the rule: $Q'(\overline{X}) \leftarrow R'_1(\overline{K}; \bar{u}_1), \ldots, R'_n(\overline{K}; \bar{u}_n)$. Let Π^{test}_{Q} be the test program containing Q, Q' and the test rule as in the previous case. Then the following holds:

$$Q \text{ is instable in } \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle \ \text{ iff } \ \Pi^{cl}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{I}} \cup \Pi^{test}_{Q} \models_{brave} \mathit{Instable}.$$

B.2 Proof of Proposition 12

In the following we prove Proposition 12 defined above.

In particular, we show how to encode the brave reasoning under Stable Model Semantics (SMS) for a given Datalog program $\Pi \cup \mathcal{D}$ with negation into stability problem for normal cyclic singleton MDBP $\langle \mathcal{P}_{\Pi,A}, \mathcal{I}_0, \mathcal{D} \rangle$ under closed semantics, where program Π is encoded in the process model and data part of the program $\mathcal{D}$ is encoded in the database of the process. As usual, the test query is Q_{test}.

Standard notation for Datalog program with negation For Datalog programs under stable model semantics (SMS) we use the following notation. A normal Datalog rule is a rule of the form

$$R(\bar{u}) \leftarrow R_1(\bar{u}_1), \ldots, R_l(\bar{u}_l), \neg R_{l+1}(\bar{u}_{l+1}), \ldots, \neg R_h(\bar{u}_h).$$

We use H to denote the head of the rule $R(\bar{u})$, and $A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h$ to denote body atoms $R_1(\bar{u}_1), \ldots, R_l(\bar{u}_l), \neg R_{l+1}(\bar{u}_{l+1}), \ldots, \neg R_h(\bar{u}_h)$ Then we can write the rules r as:

$$H \leftarrow A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h.$$

We represent a fact $R(\bar{u})$ as a Datalog *fact rule* $R(\bar{u}) \leftarrow$.

A Datalog program with negation Π is a finite set of normal Datalog rules $\{r_1, \ldots, r_k\}$.

Grounding of a Datalog program Let r be a normal Datalog rule and C a set of constants. The *grounding* $gnd_C(r)$ of r is a set of rules without variables obtained by substituting the variables in r with constants from C in all possible ways. In this way we can obtain several grounded rules from a non-grounded rule. The grounding $gnd(\Pi)$ for a program Π is a program obtained by grounding rules in Π using the constants from Π. We note that program $gnd(\Pi)$ and Π have the same semantic properties (they have the same SM, see later). Program $gnd(\Pi)$ is just an expanded version of Π (it can be exponentially bigger than Π).

Stable model semantics Concerning stable model semantics we use the following notation. An *interpretation* of a program represented as a set of facts. Let M be an interpretation. We define the *reduct* of Π for M as the ground positive program

$$\Pi^M = \{A \leftarrow A_1, \ldots, A_l \mid A \leftarrow A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h \in gnd(\Pi), \\ M \cap \{A_{l+1}, \ldots, A_h\} = \emptyset\}$$

Since Π^M is a positive ground program it has a unique Minimal Model (MM), in the inclusion sense Then,

$$M \text{ is a } \textit{stable model} \text{ (SM) of } \Pi \quad \text{iff} \quad M \text{ is the minimal model of } \Pi^M.$$

Given a program Π and a fact A we say that

$$\Pi \models_{brave} A$$

if there exists a SM M of Π such that $A \in M$.

For a given Π and a fact A, deciding whether $\Pi \models_{brave} A$ is NExpTime-hard.

Encoding of Brave Entailment into Stability Problem Given a program $\Pi \cup \mathcal{D}$ and a fact A we construct an MDBP $\mathcal{B}_{\mathcal{P},\mathcal{D},A}\langle \mathcal{P}_{\Pi,A}, \mathcal{I}_0, \mathcal{D}\rangle$ such that for a test query $Q_{test} \leftarrow dummy$ the following holds:

$$\Pi \cup \mathcal{D} \models_{brave} A \text{ iff } Q_{test} \text{ is stable in } \langle \mathcal{P}_{\Pi,A}, \mathcal{I}_0, \mathcal{D}\rangle.$$

For convenience, in the following we use Π to denote $\Pi \cup \mathcal{D}$, unless otherwise is stated.

Intuitively, process $\mathcal{B}_{\mathcal{P},\mathcal{D},A}$ is constructed such that the following holds.

- The process generates all possible interpretations for Π using the variables and constants from Π. That is, it generates all possible candidates for SMs of Π.
- For every such SM candidate M, the process checks if M is a SM of Π by:
  i) computing the MM of Π^M denoted with M';
  ii) checking if $M' = M$.
- If M is a SM of Π then the process checks for the given fact A whether it holds that $A \in M$. If so, the process produces *dummy*.

We organize $\mathcal{B}_{\Pi,A}$ in 6 subprocesses represented in Figure 1.

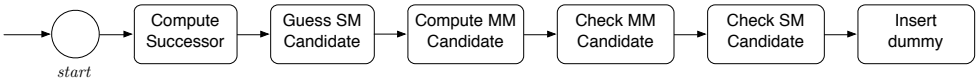

Figure 1 Subprocesses composing the process net of $\mathcal{P}_{\Pi,A}$.

The subprocesses are intuitively defined as follows:

Subp 1. *(Compute successor relations)* First we compute the successor relations $Succ^i$ of sufficient size i, that we need in the next steps. This we need for technical reasons.

Subp 2. *(Guess a SM candidate)* At this step, the process produces a SM candidate by non-deterministically producing facts obtained from relations and constants that appear in the program. Let R be a relation in Π. Then, for each R-fact that can be obtained by taking the constants from Π, a process does an execution step at *the choice place* from which if an instance traverse one way the process produces this R-facts, and if it traverses the other way then it does not.

We denote with M the guessed SM candidate.

Subp 3. *(Compute a MM candidate of the reduct)* We want compute the MM of the reduct Π^M. To do so, we first compute a candidate M' for the MM by by non-deterministically applying the rules of Π^M. Computing a candidate and the testing if the candidate is the MM is our approach to find the MM.

Subp 4. *(Check if M' is the MM of the reduct)* At this step we check if M' is indeed the minimal model of the reduct Π^M. If this is not the case, the process is not going to progress further.

Subp 5. *(Check if SM candidate is a SM)* If we are at this step then M' is the MM of Π^M. Now we check if $M' = M$. If this is the case, then M is a stable model of Π.

Subp 6. *(Insert dummy)* Finally, we check if $A \in M$. If this is the case then the process produces *dummy*.

Instance and data part. We initialize the instance part $\mathcal{I}_0$ by placing a single instance at the start place, we set database to be the data part of the program $\mathcal{D}$.

Process model. In the following we construct the process model $\mathcal{P}_{\Pi,A}$.

Subp 1: Computing successor relations. In order to nondeterministically select which R-facts to produce for a relation R in Π, we introduce sufficiently big linear order that index all R-facts. Since there are exponentially many R-facts we define the process rules that compute the order starting from an order of a polynomial size. The rules that compute the exponentially big order uses the same rules define as in Lemma 20. Here, the difference is that we use constants from Π as digits.

Let $C = \{b_1, \ldots, b_c\}$ be the constants from Π. We define a linear order $<$ on C such that

$$b_1 < b_2 < \cdots < b_c.$$

Let $<^j$ be the lexicographical order linear order on C^j, defined from $<$ for some $j > 0$. Further, let n be the maximum between

- the maximal arity of a relation in Π; and
- the largest number of variables in a rule in Π.

We want to compute the successor relation $Succ^j$ that contains immediate successors in the order $<^j$ for $j = 1, \ldots, n$

Vocabulary and Symbols To encode the order as database relations we introduce relations: *Const* of size 1 to store constants from Π; $Succ^j$ of size $2j$ to store immediate successors in the order $<^j$; $First^j$ and $Last^j$ to store the first and the last element of the order $<^j$. That is,

- $Const(b)$ – is true iff b is a constant from Π.
- $Succ^j(\bar{b}, \bar{b}')$ – is true iff $\bar{b}$ is the immediate successor of $\bar{b}'$ in the order $<^j$;
- $First^j(\bar{b})$ – is true iff $\bar{b}$ is the first element in the order $<^j$;
- $Last^j(\bar{b})$ – is true iff $\bar{b}$ is the last element in the order $<^j$.

Initialization We initialize relations for the ordering as follows:

- $Const(b)$ – we intialize relation $Const$ with all the constants from Π;
- $Succ^1(b, b')$ – we initialize relation $Succ^1$ saying that b' is the successor of b;
- $First^j(b, \ldots, b)$ – is the initialization for relation $First^j$ such that b is the first element in the order $<$;
- $Last^j(b, \ldots, b)$ – is the initialization for relation $Last^j$ such that b is the last element in the order $<$.

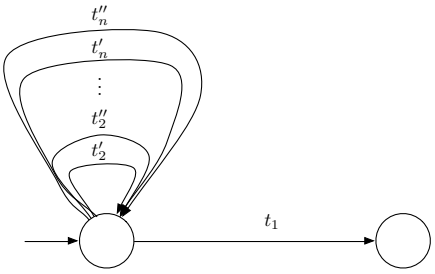

Figure 2 Subprocess 1 computes successor relations $Succ^i$ for $i = 1, \ldots, n$

Encoding into the process We introduce $2n - 1$ transitions $t_1, t'_2, t''_2, \ldots, t'_n, t''_n$ (see Figure 2) such that t'_j and t''_j are used to generate $Succ^j$. Then, we set the execution condition for these transitions to be always executable:

$$E_{t'_j} = E_{t''_j} = true.$$

We use the writing rules to populate the relations $Succ^j$ for $1 < j \leq n$ as follows:

$$W_{t'_{j+1}} : Succ^{j+1}(Z, \overline{X}, Z, \overline{Y}) \leftarrow Const(Z), Succ^j(\overline{X}, \overline{Y});$$
$$W_{t''_{j+1}} : Succ^{j+1}(Z_1, \overline{X}, Z_2, \overline{Y}) \leftarrow Succ^1(Z_1, Z_2), First^j(\overline{X}), Last^j(\overline{Y}).$$

Once all successor relations are generated transition t_1 can be executed:

$$E_{t_1} : Succ^n(\overline{X}, _), Last^n(\overline{X}).$$

Subp 2: Guessing a SM candidate. Let $R_1, \ldots, R_m$ be the relations in Π. For every relation R in Π we create a subprocess *Guess-R* that non-deterministically guesses R-facts that belong to a SM candidate M.

Subprocess 2 is composed by connecting subprocess *Guess-R* for each relation R as depicted in Figure 3.

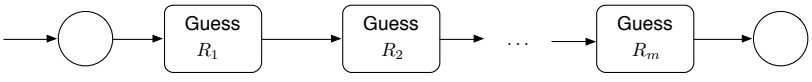

Figure 3 Subprocess 2

Notation We assume the following notations: a arity of a relation R in Π; m is the number of relations in Π; $Done_R$ is a relation of arity a such that $Done_R(\bar{u})$ is true after the subprocess *Guess-R* has guessed whether to include $R(\bar{u})$-fact in the SM candidate or not

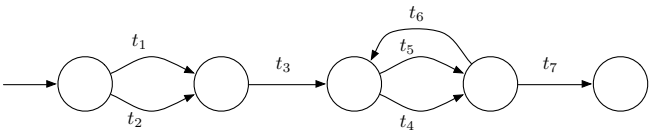

Figure 4 Subprocess *Guess-R*

Encoding into the process The subprocess *Guess-R* is defined as in Figure 4

For convenience introduce condition $Current_R(\overline{X})$ that is true if the next $R(\overline{X})$-fact for which the process has to decide whether to include it in the SM candidate or not. The condition is defined with:

$$Current_R(\overline{X}) : Succ^a(\overline{X}, \overline{Y}), Done_R(\overline{Y}), \neg Done_R(\overline{X}).$$

Transitions t_1 and t_2 are executed non-deterministically. Intuitively, they non-deterministically decide whether the R-fact, obtained by grounding R with constants from $First^a$, belongs to the SM candidate (t_1) or not (t_2):

$$\begin{aligned} &E_{t_1} = E_{t_2} : true; \\ &W_{t_1} : R(\overline{X}) \leftarrow First^a(\overline{X}); \\ &W_{t_2} : true \leftarrow true. \end{aligned}$$

Then, transition t_3 inserts that the guess for the first R-fact has been made by inserting $Done_R(\bar{x})$:

$$\begin{aligned} &E_{t_3} : true; \\ &W_{t_3} : Done_R(\overline{X}) \leftarrow First^a(\overline{X}). \end{aligned}$$

Transitions t_4 and t_5, similarly to transitions t_1 and t_2, non-deterministically guess whether the next $R(\overline{X})$-fact belongs to the SM candidate or not:

$$\begin{aligned} &E_{t_4} = E_{t_5} : true; \\ &W_{t_4} : R(\overline{X}) \leftarrow Current_R(\overline{X}); \\ &W_{t_5} : true \leftarrow true. \end{aligned}$$

Transition t_6, similarly to transition t_3, inserts fact $Done_R(\overline{X})$ after decision for $R(\overline{X})$-fact has been made:

$$\begin{aligned} &E_{t_6} : true; \\ &W_{t_6} : Done_R(\overline{X}) \leftarrow Current_R(\overline{X}). \end{aligned}$$

When all guesses have been made, transition t_7 can be executed and the next subprocess will be executed:

$$\begin{aligned} &E_{t_7} : Done_R(\overline{X}), Last^a(\overline{X}); \\ &W_{t_7} : true \leftarrow true. \end{aligned}$$

Subp 3: Compute the minimal model of the reduct. The subprocesses 3 and 4 compute the MM M' of Π^M. Intuitively, this done in the following way:

- Since Π^M is a positive ground program the MM of Π^M is unique and it can be computed as the Least Fixed Point (LFP) on the rules in Π^M.
- In subprocess 3, depicted in Figure 5, the process produces facts that are in the LFP of Π^M. For every relation R we introduce a relation R' that stores facts produced by the LFP computation.
- In principle, subprocess 3 can produce all facts from the LFP if it executes a sufficient number of times. However, it can produce also only a part of the LFP if it decides to traverse t_{k+1}.
- In other words, subprocess 3 non-deterministically decides how many facts from the LFP to produce.
- In subprocess 4 we check if all facts from the LFP of Π^M are indeed produced at subprocess 3.

Vocabulary and Symbols

- $R'(\bar{u})$ – holds iff $R(\bar{u})$ is in the LFP of Π^M (i.e. it is in the MM of Π^M) and it is computed by subprocess 3.

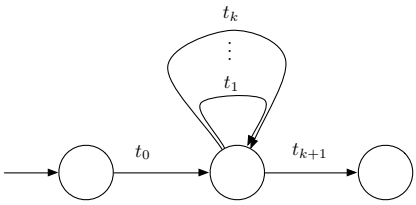

Figure 5 Subprocess 3 computes the MM candidate of the reduct

Encoding into the process Let $\{r_1, \ldots, r_k\}$ be the rules in Π. For every rule r_i of the form $H \leftarrow A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h$ we introduce transition t_i as depicted in Figure 5 with execution condition:

$$E_{t_i} : true$$

and writing rule as follows:

$$W_{t_i} : H' \leftarrow A'_1, \ldots, A'_l, A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h.$$

Here, atoms $H', A'_1, \ldots, A'_l$ are the same as $H, A_1, \ldots, A_l$, except that each relation name R is renamed with R'. Atoms $A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h$ evaluates over M and they are true iff there exists a grounding substitution θ (a substitution that replaces variables with constants) such that the ground rule $\theta A \leftarrow \theta A_1, \ldots, \theta A_l$ is in the reduct Π^M. For $l = 0$, the fact $\theta H'$ is produced by the process since the rule $\theta H \leftarrow$ is in Π^M as thus H is in the LFP of Π^M. For $l > 0$, assume that $\theta A'_1, \ldots, \theta A'_l$ are already produced by the process such that $\theta A_1, \ldots, \theta A_l$ are in the LFP of Π^M. Then we have that $\theta H'$ is produced by the process iff θH is in the LFP of Π^M.

Subp 4: Check if the computed model is a minimal model of the reduct. After the execution of subprocess 3 we obtain a MM candidate M' as a set of R'-facts produced by the process.

In this step we check if M' is indeed a MM of Π^M, because in the preceding step it may be that the process has generated a part of the LFP of Π^M.

For the check we define the process as in Figure 6, where each transition t_{r_i} checks if M' contains all facts in the LFP that can be produced by the rule r_i.

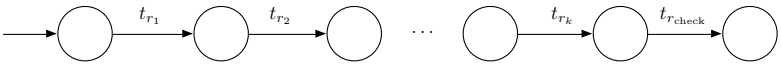

Figure 6 Subprocess 4 checks if the MM candidate is the MM of the reduct

Notation

We introduce unary predicate $fail_{MM}$ that is true iff M' is not a MM.

Encoding into the process For every rule r we introduce a transition t_r with execution condition

$$E_{t_r} : true,$$

and with writing rule as follows:

$$W_{t_r} : fail_{MM} \leftarrow A'_1, \ldots, A'_l, \neg H', A_1, \ldots, A_l, \neg A_{l+1}, \ldots, \neg A_h$$

Fact $fail_{MM}$ is produced by the process iff facts $\theta A'_1, \ldots, \theta A'_l$ are produced by the subprocess 3 while $\theta H'$ is not, for some substitution θ. Obviously, this is true iff M' is not the MM of the reduct.

Last transition t_{check} is executable if none of the previous steps has generated the $fail_{MM}$ predicate:

$$E_{t_{check}} : \neg fail_{MM}.$$

Subp 5: Checking if SM candidate is a SM. If the process execution can reach subprocess 5 it means that M' is indeed the MM of reduct Π^M. It remains to check if M is a SM of Π, that is if $M' = M$.

For this check we define the subprocess as in Figure 7.

Transition t'_i checks if there is a R'_i-fact for which there is no R_i-fact and transition t''_i checks if there is a R_i-fact for which there is no R'_i-fact.

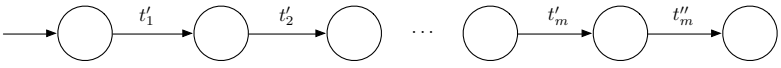

Figure 7 Subprocess 5 checks if SM candidate is a SM

Notation We introduce unary predicate $fail_i$ that holds if $M' \neq M$.

Encoding into the process Transition t'_i is encoded as follows:

$$W_{t'_i} : fail_i \leftarrow R'_i(X), \neg R_i(X).$$

Transition t''_i is encoded as follows:

$$W_{t''_i} : fail_i \leftarrow R_i(X), \neg R'_i(X).$$

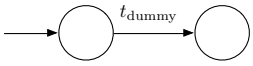

Figure 8 Subprocess 6 inserts *dummy*

Subp 6: Insert *dummy*. After the execution of subprocess 5 if no $fail_i$ facts were produced, then $M' = M$

Subprocess 6 checks whether this is the case. If $M' = M$ and $A \in M$ the process inserts *dummy*

Encoding into the process The subprocess is depicted in Figure 8 Transition t_{dummy} checks if $M' = M$ with the execution condition:

$$E_{t_{dummy}} : \neg fail_1, \ldots, \neg fail_m.$$

By traversing t_{dummy} if condition $A \in M$ then *dummy* is inserted with the following writing rule:

$$W_{t_{dummy}} : dummy \leftarrow A.$$

All together, we have that fact A is produced by the process iff there exists a SM of the program that contains A. This concludes the proof.

C Rowo MDBPs

C.1 Proof of Theorem 18

We first show encodings for fresh rowo MDBPs. Then we show encodings for closed rowo MDBPs, and finally we combine these two encodings to obtain encodings for arbitrary rowo MDBPs.

Rowo Fresh MDBPs First we analyze a fresh rowo $\mathcal{B} = \langle \mathcal{P}, \mathcal{D} \rangle$. For this case we adapt $\Pi^{po,fr}_{\mathcal{P}}$ from Section 8.

In rowo MDBPs, we have that each instance needs not to traverse a transition more than once in order to produce the most that the transition can produce. It may need to traverse some transitions more than once to reach other transitions, but in total it is sufficient that it makes at most m^2 traversals to reach all transitions, where m is the number of transitions. I.e., it is sufficient to consider executions of a single instance of maximal size m^2. Therefore, we can eliminate recursion from traversal rules in positive fresh MDBPs by creating a bounded derivation of maximal size m^2.

To this end, instead of In_p we introduce relations In^i_p for each i up to m^2 to record that a fresh instance can reach place p in i steps. That is, $In^i_p(\bar{s})$ is true iff a fresh instance with $In(\bar{s})$-record can reach place p in i steps.

Next we adapt the rules.

Traversal Rules For each transition t from a place q to a place p and for each i up to m^2 we adapt a *traversal rule* as follows:

$$In^{i+1}_p(\bar{s}) \leftarrow E^i_t(\bar{s})$$

Note that, as a difference from the general case, here $E^i_t(\bar{s})$ denotes the execution condition evaluated over the initial database (rather than on the extended database as $E'_t(\bar{s})$ would denote), and where $In(\bar{s})$ condition is replaced with $In^i_q(\bar{s})$. The database relations are not changed.

Generation Rules Similarly, for each transition t from above we introduce the following *generation rule*:

$$R'(\bar{u}) \leftarrow E^i_t(\bar{s}), B^i_t(\bar{s}', \bar{u}), \bar{s} = \bar{s}'$$

As pointed in the observations negation in E_t and B_t does not make reasoning more complex since negation is on the base relations that are not updated by the process.

Summary Let $\Pi^{ro,fr}_{\mathcal{P}}$ be the non-recursive Datalog program with stratified negation that encodes the rowo process $\mathcal{P}$ obtained from $\Pi^{po,fr}_{\mathcal{P}}$ substituting the traversal and generation rules with the rules above. The rest of the program is the same as in the positive variant.

▶ **Lemma 22.** *Let $R'(\bar{u})$ be a fact defined over adom*, then the following is equivalent:*

- *there is an execution in $\mathcal{B}$ that produces $R(\bar{u})$*
- $\Pi^{ro,fr}_{\mathcal{P}} \cup \mathcal{D} \models R'(\bar{u})$

Let Π^{test}_Q be the test program based on a query Q as defined for the general variant.

▶ **Theorem 23.** *The following is equivalent:*

- *Q is instable in $\mathcal{B}$;*
- $\Pi^{ro,fr}_{\mathcal{P}} \cup \mathcal{D} \cup \Pi^{test}_Q \models \mathit{Instable}$

Rowo Closed MDBPs We now consider a possibly cyclic rowo $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ under closed executions. We adapt the encoding from the acyclic closed variant (which can be obtained from closed positive). The main difference is that each instance is encoded independently of the others. I.e., we encode an execution of a single instance as a tuple $\bar{\omega}$ of the form

$$\bar{\omega} = \langle o; t_{h_1}, \ldots, t_{h_i} \rangle.$$

meaning that instance o traverses first t_{h_1} then t_{h_2}, and so on.

Similarly we adapt R^i's and $State^i$'s from the general case such that:

- $R^i(o; t_{h_1}, \ldots, t_{h_i}; \bar{s})$ denotes that the instance o after traversing $t_{h_1}, \ldots, t_{h_i}$ produces $R(\bar{s})$; and
- $State^i(o; t_1, \ldots, t_i; p)$ denotes that the instance o after traversing $t_1, \ldots, t_i$ is located at place p.

Similarly to the previous variant, cycles can be dealt with bounded derivations of maximal length m^2, so i ranges from $1, \ldots, m^2$. Then similarly to the closed variant, we use $In^0(o; \bar{s})$ to associate instance o with the input record $In(\bar{s})$. In this way, we obtain the facts that can be produced by each instance. Then we introduce additional rules that combine facts produced by different instances. Assume we are given a query $Q(\overline{X}) \leftarrow R_1(\bar{u}_1), \ldots, R_n(\bar{u}_n)$ that we want to check for stability. To this end, we introduce the following relations.

- *Path* is a relation with arity $m^2 + 1$ that contains legal paths of an instance. $Path(o; \bar{t}, \bar{\epsilon})$ is true iff $\bar{t}$ is a legal path in $\mathcal{P}$ for instance o. For technical reasons we introduce ϵ to denote an empty transition. Then, $\bar{\epsilon}$ is vector of ϵ that we use to fill in remaining positions in *Path* ($|\bar{\epsilon}| = m - |\bar{t}|$).
- R' is an auxiliary relation of size $1 + m^2 + arity(R)$ that we introduce for each R in $\mathcal{B}$ to store R-facts produced by an instance. That is, $R'(o; \bar{t}, \bar{\epsilon}; \bar{s})$ in true iff $R(\bar{s})$ is produced after o traversed $\bar{t}$.
- $Exec^j$ are relations of arity $(m^2 \times j) + j$ for every $j = 1, \ldots, k$ that combines legal paths for different n instances where n is is the number of atoms in the query. Then, $Exec^j(o_1, \bar{t}_1, \ldots, o_j, \bar{t}_j)$ is true iff tuple $\bar{t}_l$ is a legal path for instance o_l and if $o_h = o_l$ then $\bar{t}_h = \bar{t}_l$. This relation we use to record all combinations of instances that can contribute to create a new query answer, thus if two facts are produced by the same instance ($o_h = o_l$) then the facts have to be produced on the same legal path ($\bar{t}_h = \bar{t}_l$).

Again i ranges from $1, \ldots, m^2$. Now we define a program that generates those relations.

Initialization Rules First we adapt the *initialization rules* for a single instance:

$$State^0(o; q) \leftarrow true \text{ iff } q \text{ is the starting place of instance } o, \text{ and}$$
$$R^0(O; \overline{Y}) \leftarrow In'(O; _), R(\overline{Y}).$$

Traversal Rules Similarly, we adapt traversal rules to be for a single instance, as follows:

$$State^{i+1}(O; \overline{T}, t; p) \leftarrow State^i(O; \overline{T}; q), E_t(O)$$

for each transition t from place q to place p and $E_t(O)$ is the same as E_t except that each atom $In(\bar{s})$ is replaced with $In^0(O, \bar{s})$ and $\overline{T}$ is a vector of different variables of size i.

Generation Rules For a transition t with writing rule $W_t \colon R(\bar{u}) \leftarrow B_t(\bar{u})$, the generation rules become:

$$R^{i+1}(O; \overline{T}, t; \bar{u}) \leftarrow State^{i+1}(O; \overline{T}, t; p), B_t(O; \overline{T}; \bar{u}).$$

Here $B_t(O; \overline{T}; \bar{u})$ is obtained analogously to $E_t(O)$.

Copy Rules Then we adapt the copy rules as follows:

$$R^{i+1}(O;\overline{T},t;\overline{U}) \leftarrow R^i(O;\overline{T};\overline{U}), \text{ and}$$
$$R'(O;\overline{T},\bar{\epsilon};\overline{U}) \leftarrow R^i(O;\overline{T};\overline{U})$$

where the size of ϵ-vector $|\bar{\epsilon}| = m^2 - i$.

Summary We denote the above program as $\Pi_{\mathcal{P}}^{ro,cl}$. Then we have that the following holds:

▶ **Lemma 24.** *Let o be an instance in $\mathcal{B}$, a list of transitions $\bar{t}$ in $\mathcal{B}$ of size i, and $R(\bar{u})$ a fact, then the following is equivalent:*

- *after o traverses $\bar{t}$ the fact $R(\bar{u})$ is produced;*
- $\Pi_{\mathcal{P}}^{ro,cl} \cup \mathcal{D} \models R^i(o;\bar{t};\bar{u})$.

Combining Rules Then we need to combine the atoms produced by the different instances, e.g.,

$$R_1^i(o_1;\bar{t}_1;\bar{s}_1)$$
$$\vdots$$
$$R_n^i(o_n;\bar{t}_n;\bar{s}_n).$$

and ensure that atoms produced by one instance are all produced following one path. To do this we need to ensure that if $o_i = o_j \Rightarrow \bar{t}_i = \bar{t}_j$. This is achieved with the *combining rules*.

First we copy all legal paths of an instance from $State^i$ into $Path$ relation:

$$Path(O;\overline{T},\bar{\epsilon}) \leftarrow State^i(O;\overline{T};_) \text{ where } |\epsilon| = m^2 - i.$$

Then we initialize $Exec^1$ with all legal paths of an instance.

$$Exec^1(O;\overline{T}) \leftarrow Path(O;\overline{T}).$$

Then we combine different paths in the following way. If instance O_l executes $\overline{T}_l$ and the same instance executes $\overline{T}_{i+1}$ then $\overline{T}_l$ and $\overline{T}_{i+1}$ must be the same. This is captured with the following rules:

$$\begin{aligned} Exec^{j+1}(O_1,\overline{T}_1,\ldots,O_l,\overline{T}_l,\ldots,O_j,\overline{T}_j,O_{j+1},\overline{T}_{j+1}) \leftarrow \\ Exec^j(O_1,\overline{T}_1,\ldots,O_l,\overline{T}_l,\ldots,O_j,\overline{T}_j), \\ Path(O_{j+1};\overline{T}_{j+1}), \\ O_l = O_{j+1}, \overline{T}_l = \overline{T}_{j+1} \end{aligned}$$

for every $l \in 1,\ldots,i$.

If the instance O_{j+1} is different from all the other instances $O_1,\ldots,O_j$ then executing path of O_{j+1} can be any legal path

$$\begin{aligned} Exec^{j+1}(O_1,\overline{T}_1,\ldots,O_l,\overline{T}_l,\ldots,O_j,\overline{T}_j,O_{j+1},\overline{T}_{j+1}) \leftarrow \\ Exec^j(O_1,\overline{T}_1,\ldots,O_l,\overline{T}_l,\ldots,O_j,\overline{T}_j), \\ Path(O_{j+1};\overline{T}_{j+1}), \\ \neg(O_1 = O_{j+1}),\neg(O_2 = O_{j+1}),\ldots,\neg(O_j = O_{j+1}). \end{aligned}$$

Q'-rule Then, Q' collects what has been produced for relations $R_1, \ldots, R_n$ for the give query $Q(\overline{X}) \leftarrow R_1(\bar{u}_1), \ldots, R_n(\bar{u}_n)$ with the rule

$$\begin{aligned} Q'(\overline{X}) \leftarrow & \mathit{Exec}^n(O_1, \overline{W}_1, \ldots, O_n, \overline{W}_n), \\ & R'_1(O_1; \overline{W}_1; \bar{u}_1), \\ & \ldots \\ & R'_n(O_n; \overline{W}_n; \bar{u}_n). \end{aligned}$$

Test Rule The *test rule* is then as before:

$$\mathit{Instable} \leftarrow Q'(\overline{X}), \neg Q(\overline{X}).$$

Summary Let us denote with $\Pi^{test,ro}_{\mathcal{P},Q}$ the testing program for Q defined above. The program is non-recursive Datalog with stratified negation.

Let $\mathcal{D}_{\mathcal{I}}$ be the database that encodes the instance part *In* and that contains database $\mathcal{D}$.

▶ **Theorem 25.** *The following are equivalent:*

- *Q is instable in $\mathcal{B}$ under closed executions;*
- $\Pi^{ro,cl}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{I}} \cup \Pi^{test,ro}_{\mathcal{P},Q} \models \mathit{Instable}$.

Rowo MDBPs For arbitrary rowo, similarly to the positive variants, the encoding is obtained combining the encodings for the fresh and the closed variants.

To combine what comes from the instances in the process and the new ones it is enough to add rules that will combine the program for closed ($\Pi^{ro,cl}_{\mathcal{P},Q}$) and fresh($\Pi^{ro,fr}_{\mathcal{P},Q}$).

To this end, for a given query $Q(\overline{X}) \leftarrow R_1(\bar{u}_1), \ldots, R_n(\bar{u}_n)$, we introduce relations:

- B^i_Q of arity $\mathit{arity}(R_1) + \cdots + \mathit{arity}(R_n)$ that contains on the first i arguments what comes from a mixture of existing and new process instances while the others come only from existing process instances, for $i = 1, \ldots, n$.

Encoding into Non-Recursive Datalog Let $\mathcal{B} = \langle \mathcal{P}, \mathcal{I}, \mathcal{D} \rangle$ be a rowo MDBP.

Now we define rules that compute relations introduced above.

First we consider what is produced by the running instances

$$\begin{aligned} B^0_Q(\overline{Y}_1, \ldots, \overline{Y}_n) \leftarrow & \mathit{Exec}^n(O_1, \overline{W}_1, \ldots, O_n, \overline{W}_n), \\ & R'_1(O_1; \overline{W}_1; \bar{u}_1), \\ & \ldots \\ & R'_n(O_n; \overline{W}_n; \bar{u}_n). \end{aligned}$$

Then, for the i-th atom we both consider the case in which it was produced by a new instance (1) and the case it was produced by the instances already in the process (2). These cases are added to the combinations obtained for the atoms from 1 to $i-1$. We do this for every $i = 1, \ldots, n$.

$$\begin{aligned} B^i_Q(\ldots, & \overline{Y}_{i-1}, \overline{Y}_i, \overline{Y}_{i+1}, \ldots) \leftarrow \\ & B^{i-1}_Q(\ldots, \overline{Y}_{i-1}, _, \overline{Y}_{i+1}, \ldots), R'_i(\overline{Y}_i) \end{aligned} \tag{4}$$

$$\begin{aligned} B^i_Q(\ldots, & \overline{Y}_{i-1}, \overline{Y}_i, \overline{Y}_{i+1}, \ldots) \leftarrow \\ & B^{i-1}_Q(\ldots, \overline{Y}_{i-1}, \overline{Y}_i, \overline{Y}_{i+1}, \ldots). \end{aligned} \tag{5}$$

Then, we add the *Q'-rule* to collect what has been produced by the process for relations $R_1(\bar{u}_1), \ldots, R_n(\bar{u}_n)$ as follows:

$$Q'(\overline{X}) \leftarrow B^n_Q(\overline{Y}_1, \ldots, \overline{Y}_n).$$

The above rules extend the testing program $\Pi^{test,ro}_{\mathcal{P},Q}$ for normal cyclic arbitrary closed. We denote the new testing program with $\Pi^{test,ro,op}_{\mathcal{P},Q}$.

- Q is instable in $\mathcal{B}$;
- $\Pi^{ro,cl}_{\mathcal{P}} \cup \Pi^{ro,fr}_{\mathcal{P}} \cup \mathcal{D}_{\mathcal{I}} \cup \Pi^{test,ro,op}_{\mathcal{P},Q} \models \mathit{Instable}$.

Then we set $\Pi^{ro}_{\mathcal{P},Q} = \Pi^{ro,cl}_{\mathcal{P}} \cup \Pi^{ro,fr}_{\mathcal{P}} \cup \Pi^{test,ro,op}_{\mathcal{P},Q}$, and the claim follows from there.

C.2 Proof of Proposition 19

Proof. Assume the instance o is at place p and it has an input record $I(\bar{s}) = M_S(o)$. To show the claim it is sufficient to guess a closed execution Υ consisting of the traversals by o, and then verify whether atoms $A_1, \ldots, A_n$ can be produced by such execution. In the case of singleton rowo MDBPs under closed executions, a closed execution is uniquely determined by a path in $\mathcal{P}$. Thus, we guess a path $t_1, \ldots, t_m$ in $\mathcal{P}$ that starts in p. This guess is polynomial in the size of $\mathcal{P}$. For all transitions on the path, we further guess assignments $\alpha_1, \ldots, \alpha_m$ for the execution conditions $E_{t_1}, \ldots, E_{t_m}$. Then for each writing rule W_{t_i} of the transition t_i we guess up to n assignments $\beta^l_{t_i}$ for $1 \leq l \leq n$, because a rule may need to produce more than one fact, but no more than the size of the query. In principle, only a subset of W_{t_1}, $\ldots$, W_{t_m} may be needed to produce atoms $A_1, \ldots, A_n$. Wlog we can guess the assignments for all. Now we verify. Firstly, we verify whether the path can be traversed by the instance. This is, if for every execution condition E_{t_i} the ground query $\alpha_i E_{t_i}$ evaluates to *true* in $\mathcal{D} \cup \{In(\bar{s})\}$. Secondly, we verify whether $A_1, \ldots, A_n$ are produced on the path by checking if for every A_i there exist a writing rule $W_{t_j} : A_{t_j} \leftarrow B_{t_j}$ and assignment $\beta^l_{t_j}$ such that the ground query $\beta^l_{t_j} B_{t_j}$ evaluates to *true* in $\mathcal{D} \cup \{In(\bar{s})\}$ and A_i is equal with the head of the writing rule $\beta^l_{t_j} A_{t_j}$. Since all guesses and checks are polynomial in the size of $\mathcal{B}$ the claim follows directly. ◀